

PATENT
EMCR:070
EMC-00-174

APPLICATION FOR UNITED STATES LETTERS PATENT
for

ADAPTIVE BIT RATE CONTROL FOR RATE REDUCTION OF MPEG CODED
VIDEO

by

Seyfullah H. Oguz
Sorin Faibish
Wayne W. Duso

EXPRESS MAIL MAILING LABEL NUMBER EL291391701 US

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to processing and storage of compressed visual data, and in particular the processing and storage of compressed visual data for bit rate reduction.

2. Background Art

It has become common practice to compress audio/visual data in order to reduce the capacity and bandwidth requirements for storage and transmission. One of the most popular audio/video compression techniques is MPEG. MPEG is an acronym for the Moving Picture Experts Group, which was set up by the International Standards Organization (ISO) to work on compression. MPEG provides a number of different variations (MPEG-1, MPEG-2, etc.) to suit different bandwidth and quality constraints. MPEG-2, for example, is especially suited to the storage and transmission of broadcast quality television programs.

For the video data, MPEG provides a high degree of compression (up to 200:1) by encoding 8 x 8 blocks of pixels into a set of discrete cosine transform (DCT) coefficients, quantizing and encoding the coefficients, and using motion compensation techniques to encode most video frames as predictions from or between other frames. In particular, the encoded MPEG video stream is comprised of a series of groups of pictures (GOPs), and each GOP begins with an independently encoded (intra) I frame and may include one or more following P frames and B frames. Each I frame can be decoded without information from any preceding and/or following frame. Decoding of a P frame requires information from a preceding frame in the GOP. Decoding of a B frame requires information from both a preceding and a following frame in the GOP. To minimize

1 decoder buffer requirements, transmission orders differ from presentation orders for some
2 frames, so that all the information of the other frames required for decoding a B frame
3 will arrive at the decoder before the B frame.

4 In addition to the motion compensation techniques for video compression, the
5 MPEG standard provides a generic framework for combining one or more elementary
6 streams of digital video and audio, as well as system data, into single or multiple program
7 transport streams (TS) which are suitable for storage or transmission. The system data
8 includes information about synchronization, random access, management of buffers to
9 prevent overflow and underflow, and time stamps for video frames and audio packetized
10 elementary stream packets embedded in video and audio elementary streams as well as
11 program description, conditional access and network related information carried in other
12 independent elementary streams. The standard specifies the organization of the
13 elementary streams and the transport streams, and imposes constraints to enable
14 synchronized decoding from the audio and video decoding buffers under various
15 conditions.

16 The MPEG-2 standard is documented in ISO/IEC International Standard (IS)
17 13818-1, "Information Technology-Generic Coding of Moving Pictures and Associated
18 Audio Information: Systems," ISO/IEC IS 13818-2, "Information Technology-Generic
19 Coding of Moving Pictures and Associated Audio Information: Video," and ISO/IEC IS
20 13818-3, "Information Technology-Generic Coding of Moving Pictures and Associated
21 Audio Information: Audio," which are incorporated herein by reference. A concise
22 introduction to MPEG is given in "A Guide to MPEG Fundamentals and Protocol

1 Analysis (Including DVB and ATSC),” Tektronix Inc., 1997, incorporated herein by
2 reference.

3 MPEG-2 provides several optional techniques that allow video coding to be
4 performed in such a way that the coded MPEG-2 stream can be decoded at more than one
5 quality simultaneously. In this context, the word “quality” refers collectively to features
6 of a video signal such as spatial resolution, frame rate, and signal-to-noise ratio (SNR)
7 with respect to the original uncompressed video signal. These optional techniques are
8 known as MPEG-2 scalability techniques. In the absence of the optional coding for such
9 a scalability technique, the coded MPEG-2 stream is said to be nonscalable. The MPEG-
10 2 scalability techniques are varieties of layered or hierarchical coding techniques, because
11 the scalable coded MPEG-2 stream includes a base layer that can be decoded to provide
12 low quality video, and one or more enhancement layers that can be decoded to provide
13 additional information that can be used to enhance the quality of the video information
14 decoded from the base layer. Such a layered coding approach is an improvement over a
15 simulcast approach in which a coded bit stream for a low quality video is transmitted
16 simultaneously with an independently coded bit stream for high quality video. The use of
17 video information decoded from the base layer for reconstructing the high quality video
18 permits the scalable coded MPEG-2 stream to have a reduced bit rate and data storage
19 requirement than a comparable simulcast data stream.

20 The MPEG-2 scalability techniques are useful for addressing a variety of
21 applications, some of which do not need the high quality video that can be decoded from
22 a nonscalable coded MPEG stream. For example, applications such as video
23 conferencing, video database browsing, and windowed video on computer workstations

1 do not need the high quality provided by a nonscalable coded MPEG-2 stream. For
 2 applications where the high quality video is not needed, the ability to receive, store, and
 3 decode an MPEG-2 base-layer stream having a reduced bit rate or data storage capacity
 4 may provide a more efficient bandwidth versus quality tradeoff, and a more efficient
 5 complexity versus quality tradeoff. A scalable coded MPEG-2 stream provides
 6 compatibility for a variety of decoders and services. For example, a reduced complexity
 7 decoder for standard television could decode a scalable coded MPEG-2 stream produced
 8 for high definition television. Moreover, the base layer can be coded for enhanced error
 9 resilience and can provide video at reduced-quality when the error rate is high enough to
 10 preclude decoding at high quality.

11 The MPEG scaling techniques are set out in sections 7.7 to 7.11 of the MPEG-2
 12 standard video encoding chapter 13818-2. They are further explained in Barry G.
 13 Haskell et al., Digital Video: An Introduction to MPEG-2, Chapter 9, entitled "MPEG-2
 14 Scalability Techniques," pp. 183-229, Chapman & Hall, International Thomson
 15 Publishing, New York, N.Y., 1997, incorporated herein by reference. The MPEG
 16 scalability techniques include four basic techniques, and a hybrid technique that
 17 combines at least two of the four basic techniques. The four basic techniques are called
 18 data partitioning, signal-to-noise ratio (SNR) scalability, spatial scalability, and temporal
 19 scalability.

20 The conventional MPEG scalability techniques permit transmission of the coded
 21 video to be switched from a high quality, high bit rate stream to a low quality, low bit rate
 22 stream when transmission of the high quality, high bit rate stream is either precluded by
 23 network congestion or is not needed at the destination of the stream. However, the

1 conventional MPEG scalability techniques do not permit the bit rate reduction to be
2 freely selected. For many applications, the required bit rate is intermediate that of the
3 high quality, high bit rate stream, and the low quality, low bit rate stream provided by the
4 conventional MPEG scalability techniques. In some applications, the required bit rate
5 will fluctuate between a high rate and a low rate. In any case, it is desired to create a
6 valid MPEG data stream with the best video quality given the required bit rate.

7 Much research has been done addressing the problem of encoding a valid MPEG
8 data stream given a required constant or variable bit rate. In general, to satisfy the MPEG
9 rate control requirements, all of the data for each picture must be within the video buffer
10 at the instant it is needed by the decoder. This requirement usually translates to upper
11 and lower bounds on the number of bits allowed in each picture. For example, a number
12 of bits are allocated to each picture based on the picture type, and the bits for each picture
13 are allocated to 8x8 blocks in each picture based on a measure of local coding complexity
14 in each picture. A quantization scale is selected for each 8x8 block to encode the 8x8
15 block with the number of bits allocated to the block. See Gonzales et al., U.S. Patent
16 5,231,484 issued July 27, 1993, on Motion Video Compression System with Adaptive Bit
17 Allocation and Quantization, and Ramamurthy et al., U.S. Patent 5,675,384, issued
18 October 7, 1997.

19 Unfortunately, when the bit rate was reduced to a small fraction of the standard
20 bit rate, the bit rate control provided by allocating bits based on picture type, and
21 selecting a quantization scale for each block, became rather imprecise. Sometimes the
22 poor bit rate control was compensated for by liberal use of stuffing to make up for a
23 difference between the allocated number of bits and the actual number of bits for each

The basic objective of the present invention is to provide an improved bit rate control for producing reduced-quality MPEG coded video from original quality MPEG coded video. It is desired to produce the reduced-quality MPEG coded video at a given constant or variable bit rate with a minimum of computing resources, yet ensure that the reduced-quality MPEG coded video will have a relatively constant visual picture quality over the picture segments and thus minimal degradation in picture quality.

In accordance with one aspect, the invention provides a method of producing reduced-quality MPEG coded video from original-quality MPEG coded video. The original-quality MPEG coded video includes a set of non-zero AC discrete cosine transform (DCT) coefficients for 8x8 blocks in frames of the original-quality MPEG coded video. The reduced-quality MPEG coded video also has frames of 8x8 blocks. Each frame in the reduced-quality MPEG coded video has a corresponding frame in the original-quality MPEG coded video, and each 8x8 block in each frame of the reduced-quality MPEG coded video has a corresponding block in a corresponding frame in the original-quality MPEG coded video. The method includes the steps of selecting non-zero AC DCT coefficients from corresponding blocks in the original-quality MPEG coded video to be included in the blocks of the reduced-quality MPEG coded video so that each 8x8 block in each frame of the reduced-quality MPEG coded video has a number of bits encoding non-zero AC DCT coefficients that is generally proportional to the number of bits encoding non-zero AC DCT coefficients for the corresponding 8x8 block in the corresponding frame of original-quality MPEG coded video.

In accordance with another aspect, the invention provides a method of producing reduced-quality MPEG coded video from original-quality MPEG coded video. The original-quality MPEG coded video includes a set of non-zero AC discrete cosine transform (DCT) coefficients for 8x8 blocks in frames of the original-quality MPEG coded video. The reduced-quality MPEG coded video also has frames of 8x8 blocks. Each frame in the reduced-quality MPEG coded video has a corresponding frame in the original-quality MPEG coded video, and each 8x8 block in each frame of the reduced-quality MPEG coded video has a corresponding block in a corresponding frame in the

1 original-quality MPEG coded video. The method includes a number of steps for each
 2 block in the reduced-quality MPEG coded video. These steps include (a) determining the
 3 number of bits used in encoding non-zero AC DCT coefficients in the corresponding
 4 block of original-quality MPEG coded video; (b) computing a number of bits available
 5 for encoding AC DCT coefficients in the original-quality MPEG coded video by scaling
 6 the number of bits used in encoding non-zero AC DCT coefficients in the corresponding
 7 block of original-quality MPEG coded video with a scale factor; and (c) selecting non-
 8 zero AC DCT coefficients in a certain order from the corresponding block in the original-
 9 quality MPEG coded video to be included in said each block of the reduced-quality
 10 MPEG coded video until the number of bits available for encoding the AC DCT
 11 coefficients in the block in the reduced-quality encoded video is not sufficient for
 12 encoding, in the block of the reduced-quality MPEG coded video, any more of the AC
 13 DCT coefficients in the corresponding block of original-quality MPEG coded video.

14 In accordance with yet another aspect, the invention provides a method of
 15 producing in real-time a stream of reduced-quality MPEG-2 coded video from a source of
 16 original-quality MPEG-2 coded video. The original-quality MPEG-2 coded video
 17 includes a set of non-zero AC discrete cosine transform (DCT) coefficients for 8x8
 18 blocks in frames of the original-quality MPEG-2 coded video. The reduced-quality
 19 MPEG-2 coded video also has frames of 8x8 blocks. Each frame in the reduced-quality
 20 MPEG-2 coded video has a corresponding frame in the original-quality MPEG-2 coded
 21 video, and each 8x8 block in each frame of the reduced-quality MPEG-2 coded video
 22 having a corresponding block in a corresponding frame in the original-quality MPEG-2
 23 coded video. The method includes a number of steps for each frame in the reduced-

quality MPEG-2 coded video. The steps for each frame include (a) computing a moving average of the size of the corresponding frame in the original-quality MPEG-2 coded video; (b) computing a scale factor from the moving average of the size of the corresponding frame in the original-quality MPEG-2 coded video and a desired size of said each frame of the reduced-quality MPEG-2 coded video; and (c) for each 8x8 block in said each frame: (i) determining the number of bits used in encoding non-zero AC DCT coefficients in the corresponding block of original-quality MPEG-2 coded video; (ii) computing a number of bits available for encoding AC DCT coefficients in the original-quality MPEG-2 coded video by scaling the number of bits used in encoding non-zero AC DCT coefficients in the corresponding block of original-quality MPEG-2 coded video with a scale factor, and (iii) selecting non-zero AC DCT coefficients in a parsing order from the corresponding block in the original-quality MPEG-2 coded video to be included in said each block of the reduced-quality MPEG-2 coded video until the number of bits available for encoding the AC DCT coefficients in the block in the reduced-quality encoded video is not sufficient for encoding, in the block of the reduced-quality MPEG-2 coded video, any more of the AC DCT coefficients in the corresponding block of original-quality MPEG-2 coded video.

BRIEF DESCRIPTION OF THE DRAWINGS

Other objects and advantages of the invention will become apparent upon reading the following detailed description with reference to the accompanying drawings, in which:

FIG. 1 is a block diagram of a data network including a video file server implementing various aspects of the present invention;

1 FIG. 2 is a flowchart of a procedure executed by a stream server computer in the
2 video file server of FIG. 1 to service client requests;

3 FIG. 3 is a flowchart of a procedure for splicing MPEG clips;

4 FIG. 4 is a flowchart of a procedure for seamless video splicing of MPEG clips;

5 FIG. 5 is a more detailed flowchart of the procedure for seamless video splicing
6 of MPEG clips;

7 FIG. 6 is a continuation of the flowchart begun in FIG. 5;

8 FIG. 7 is a timing diagram showing a timing relationship between video
9 presentation units (VPUs) and associated audio presentation units (APUs) in an original
10 MPEG-2 coded data stream;

11 FIG. 8 is a timing diagram showing a timing relationship between video
12 presentation units (VPUs) and associated audio presentation units (APUs) for a fast-
13 forward trick-mode stream;

14 FIG. 9 is a flowchart of a procedure for selection and alignment of audio
15 presentation units (APUs) in the fast-forward trick-mode stream;

16 FIG. 10 is a flowchart of a procedure for producing a trick-mode MPEG-2
17 transport stream from a regular MPEG-2 transport stream (TS);

18 FIG. 11 is a diagram illustrating relationships between the MPEG discrete cosine
19 transform (DCT) coefficients, spatial frequency, and the typical zig-zag scan order;

20 FIG. 12 is a diagram illustrating a relationship between an MPEG-2 coded bit
21 stream and a reduced-quality MPEG-2 coded bit stream resulting from truncation of high-
22 order DCT coefficients;

FIG. 13 is a flowchart of a procedure for scaling MPEG-2 coded video using a variety of techniques;

FIG. 14 is a flowchart of a procedure for signal-to-noise ratio scaling MPEG-2 coded video using a frequency-domain low-pass truncation (FDSNR_LP) technique;

FIG. 15 is a flowchart of a procedure for signal-to-noise ratio scaling MPEG-2 coded video using a frequency-domain largest-magnitude coefficient selection (FDSNR_LM) technique;

FIG. 16 is a flowchart of a procedure that selects one of a number of techniques for finding a certain number "k" of largest values out of a set of "n" values;

FIG. 17 is a flowchart of a procedure for finding a certain number "k" of largest values from a set of "n" values, which is used in the procedure of FIG. 16 for the case of $k \ll \frac{1}{2} n$;

FIG. 18 is a diagram of a hash table and associated hash lists;

FIG. 19 is a flowchart of a procedure for finding a certain number "k" of values that are not less than the smallest of the "k" largest values in a set of "n" values beyond a certain amount.

FIG. 20 is a flowchart of modification of the procedure of FIG. 15 in order to possibly eliminate escape sequences in the (run, level) coding of the largest magnitude coefficients;

FIG. 21 is a flowchart of a subroutine called in the flowchart of FIG. 20 in order to possibly eliminate an escape sequence;

FIG. 22 is a first portion of a flowchart of a procedure for scaling an MPEG-2 coded video data stream using the modified procedure of FIG. 20 while adjusting the

parameter “k” to achieve a desired bit rate, and adjusting a quantization scaling factor (QSF) to achieve a desired frequency of occurrence of escape sequences;

FIG. 23 is a second portion of the flowchart begun in FIG. 22;

FIG. 24 is a simplified block diagram of a volume containing a main file, a corresponding fast forward file for trick mode operation, and a corresponding fast reverse file for trick mode operation;

FIG. 25 is a more detailed block diagram of the volume introduced in FIG. 24;

FIG. 26A is a diagram showing video file access during a sequence of video operations including transitions between the main file, the related fast forward file, and the related fast reverse file;

FIG. 26B shows a script of a video command sequence producing the sequence of video play shown in FIG. 26A;

FIG. 27 is a table of read and write access operations upon the volume of FIG. 24 and access modes that are used for the read and write access operations;

FIG. 28 is a hierarchy of video service classes associated with the fast forward file and the fast reverse file in the volume of FIG. 25;

FIG. 29 shows a system for modifying and combining an MPEG-2 audio-visual transport stream with an MPEG-2 closed-captioning transport stream to produce a multiplexed MPEG-2 transport stream having the same bit rate as the original MPEG-2 audio-visual transport stream;

FIG. 30 is a diagram showing blocks in a frame of original-quality MPEG coded video being proportionally reduced to obtain a frame of reduced-quality MPEG coded video in an adaptive bit rate reduction method;

1 FIG. 31 shows a graph of frame size in bits per frame as a function of frame
2 number for original-quality MPEG coded video and corresponding reduced-quality
3 MPEG coded video obtained using the adaptive bit rate reduction method introduced in
4 FIG. 30;

5 FIG. 32 shows a graph of frame size in bits per frame as a function of frame
6 number for original-sized MPEG I frames and corresponding reduced-size I frames for
7 trick file generation using the adaptive bit rate reduction method introduced in FIG. 30;

8 FIG. 33 is a first sheet of a flowchart for programming a digital computer to
9 implement the adaptive bit rate reduction method introduced in FIG. 30;

10 FIG. 34 is a second sheet of the flowchart begun in FIG. 33;

11 FIG. 35 is a third sheet of the flowchart begun in FIG. 33;

12 FIG. 36 is a subroutine for programming a digital computer to determine a
13 coefficient bit rate reduction factor for a reduction from an MPEG source having a known
14 constant bit rate;

15 FIG. 37 is a subroutine for programming a digital computer to determine a
16 coefficient bit rate reduction factor for a reduction from an MPEG source having an
17 unknown or variable bit rate; and

18 FIG. 38 is a block diagram illustrating the use of the adaptive bit rate reduction
19 method introduced in FIG. 30 for the real-time production of a reduced-quality MPEG-2
20 stream from an original-quality MPEG-2 source.

21 While the invention is susceptible to various modifications and alternative forms,
22 specific embodiments thereof have been shown by way of example in the drawings and
23 will be described in detail. It should be understood, however, that it is not intended to

limit the form of the invention to the particular forms shown, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the scope of the invention as defined by the appended claims.

DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

With reference to FIG. 1, there is shown a block diagram of a data network 20 linking a number of clients 21, 22, 23 to a video file server 24 implementing various aspects of the present invention. The video file server 24 includes at least one stream server computer 25 and a data storage system 26. The stream server computer 25 has a processor 27 and a network link adapter 28 interfacing the processor to the data network 20. The processor 27 executes a data streaming program 29 in memory 30 in order to stream MPEG coded video in real-time to the clients.

Client requests for real-time video are placed in client play lists 31 in order to schedule in advance video file server resources for the real-time streaming of the MPEG coded video. The play lists 31 specify a sequence of video clips, which are segments of MPEG-2 files 32, 33 in data storage 34 of the data storage system 26. The stream server processor 27 accesses a client play list in advance of the time to begin streaming MPEG coded video from a clip, and sends a video prefetch command to a storage controller 35 in the data storage system 26. The storage controller responds to the video prefetch command by accessing the clip in the data storage 34 to transfer a segment of the clip to cache memory 36. When the video data of the segment needs to be sent to the client, the stream server processor 27 requests the data from the storage controller 35, and the storage controller immediately provides the video data from the cache memory 36. Further details regarding a preferred construction and programming of the video file

1 server 24 are disclosed in Duso et al., U.S. Patent 5,892,915 issued Apr. 6, 1999, entitled
2 "System Having Client Sending Edit Commands to Server During Transmission Of
3 Continuous Media From One Clip in Play List for Editing the Play List," incorporated
4 herein by reference.

5 In accordance with an aspect of the invention, the stream server computer 25
6 executes an MPEG scaling program 38 to produce reduced-quality MPEG coded video
7 from nonscalable MPEG-2 coded video by truncating discrete cosine transform (DCT)
8 AC coefficients from the coded blocks in the MPEG-2 coded video data. The reduced-
9 quality MPEG coded video can be produced during ingestion of an MPEG-2 file 32 from
10 the network 20, and stored in one or more associated files 37. Alternatively, the reduced-
11 quality MPEG coded video in the files 37 could be produced as a background task from
12 the MPEG-2 file 32. Reduced-quality MPEG coded video could also be produced in
13 real-time from an MPEG-2 file 33 during streaming of the reduced-quality MPEG coded
14 video from the stream server computer 25 to the network 20. The reduced-quality MPEG
15 coded video is useful for a variety of applications, such as browsing and review of stored
16 MPEG-2 assets for search and play-list generation, bit stream scaling for splicing, and
17 bit-rate adjustment via video quality alteration for services with limited resources.

18 A typical example of browsing for play-list generation involves searching stored
19 assets in a multi-media data base for segments of a desired content to be included in the
20 play list, and in particular selecting the beginning frame and ending frame of each
21 segment to be included. Such editing occurs often in the broadcast environment for
22 inserting commercials and news clips into pre-recorded television programming, and for
23 editing movies for content and time compression. The decoding technique of the present

1 invention permits a PC workstation 23 to perform the decoding and display in real-time
2 by execution of a software program. An operator can view the video content in a display
3 window 39 in a fast-forward or fast-reverse mode, stop at and resume from freeze frames
4 that are valid "in points" and "out points" for seamless splicing, and select an in-point
5 and out-point for a next segment to be included in the play list. The stream server
6 computer 25 could also include a seamless splicing program 40 providing seamless
7 transitions between video segments that are contiguous in a play list and are from
8 different video clips.

9 For seamless splicing, it is often necessary to reduce the bitrate for one or more
10 frames at the end of a first segment prior to splicing to a second segment. In this case the
11 bitrate must be reduced to avoid buffer overflow as a result of displaying the original
12 frames at the end of the first segment. One method of reducing the bitrate is to insert a
13 freeze frame at the end of the first segment, but this has the disadvantage of introducing
14 distortion in the temporal presentation of the frames and precluding frame accuracy. A
15 less disruptive method is to use the present invention for reducing the bitrate for a lower-
16 quality presentation of one or more frames at the end of the first segment.

17 The present invention can also reduce the bit transmission rate and storage
18 requirements for MPEG-2 applications by altering the video quality. For example,
19 different clients may present different bandwidth access requests for video from
20 nonscalable MPEG-2 files 32, 33 in the video file server. Also, temporary network
21 congestion may limit the bandwidth available to satisfy a request for real-time streaming
22 of video data. In each case, the present invention can alter the video quality to meet the
23 desired or available bandwidth to satisfy the request.

1 to the fast forward and fast reverse playback functions of a video cassette recorder
2 (VCR). The problem with trick-mode operation is that the speed of the MPEG stream
3 cannot simply be speeded up because the transmission bandwidth would be excessive and
4 a conventional MPEG-2 decoder will not be able to handle the increased data rate or even
5 if the decoder would have been able to support the increased data rate, such a change in
6 the original operating conditions is not allowable. For this reason, in trick-mode, neither
7 the original display rate of 29.97 frames per second (for NTSC or 25 frames per second
8 for PAL) nor the original transport stream (TS) multiplex rate should change. Nor is it
9 possible to simply decimate frames since only the I frames are independently coded, and
10 the P frames and B frames need the content of certain other frames for proper decoding.
11 The I frames typically occur once for every 15 frames. Assuming that this convention is
12 followed in the encoding process, it would be possible to preserve and play each I frame
13 from each and every group of pictures (GOP), resulting in a 15 times slower temporal
14 sampling rate, or a 1 to 15 speeding up of motion if the I frames only are played back at
15 the nominal NTSC rate of approximately 30 frames per second. Consequently, the
16 content of a 60 minutes duration clip will be covered in 4 minutes. Unfortunately the
17 average information content per frame for the I frames is more than four times the
18 average information content of the P and B frames. Therefore, the trick-mode cannot be
19 implemented simply by transmitting only the I frames for a speed-up by a factor of 15,
20 because this would need an increase in the TS multiplex rate over the nominal rate.

21 In particular, the average information content of an I frame has been measured to
22 be about 56,374.6 bytes. If the I frames only are transmitted at the standard NTSC rate,
23 then the bit transmission rate would be: $8(\text{bits per byte}) * 56,374.6(\text{bytes per frame}) *$

1 29.97(frames per sec.) or about 13,516,374.1 bits per second only for the video stream,
2 which is significantly above - almost 3.38 times - the original rate of 4 megabits per
3 second used in this test. This calculation, being based on an average quantity, is ignoring
4 the indispensable need for an actually higher transport rate to provide some safety margin
5 to handle short-term-sustained large size I frame chains (bursts) which practically always
6 happen. Clearly, some form of modification in the trick-mode operation definition is
7 required to handle this problem and pull the bit-rate requirement down to the nominal 4
8 megabits per second.

9 Two degrees of freedom are available to achieve such a reduction in the required
10 bit-rate for trick-mode operation. The first is I frame compression quality and the second
11 is a motion speed-up ratio. With respect to compression quality, it is well known that
12 human observers' perception of image detail degrades with increasing motion speed of
13 objects in the scene. Based on this fact, the type of D pictures were introduced in MPEG-
14 1 video syntax for fast visible (forward or reverse) search purposes. (See ISO/IEC 11172-
15 2: 1993 Information Technology - Coding of moving pictures and associated audio for
16 digital storage media at up to about 1.5 Mbits/s - Part 2: Video, Annex D.6.6. Coding D-
17 Pictures, p. 102). D pictures make use of only the DC coefficients in intra coding to
18 produce very low quality (in terms of SNR) reproductions of desired frames which were
19 judged to be of adequate quality in fast search mode.

20 In order to provide support for enhanced quality trick-mode operation, the quality
21 of the original I frames can be reduced by the preservation of just a sufficient number of
22 AC DCT coefficients to meet the bit-rate limitation. Based on experiments with two
23 standard video test sequences (one encoded at 15 Mbits/sec. and the other at 24

1 Mb/s/sec. and both with I frames only), it is observed that the bandwidth for I frames can
2 be scaled to one half by keeping about 9 lowest order AC coefficients and eliminating the
3 rest. This scheme provides good quality even at the full spatial and temporal resolution,
4 much better than D pictures.

5 The inherent speed-up ratio lower bound imposed by the GOP structure can be
6 relaxed and further lowered by freeze (P) frame substitution in between genuine (SNR
7 scaled or non-scaled) I frames. The maximum number of freeze frames that can be
8 inserted before visually disturbing motion jerkiness occurs, is very likely to depend
9 heavily on the original GOP structure (equivalently the separation between I frames of
10 the original sequence) and the original amount of motion in the clip. However, 1, 2 or 3
11 freeze frame substitutions in between genuine I frames present reasonable choices which
12 will yield speed-up ratios of 1 to 7.5, 1 to 5 and 1 to 3.75 respectively instead of the 1 to
13 15 speed-up ratio provided by the genuine I frames only implementation. (These ratios
14 are computed by a first-order approximation that neglects a slight increase in bandwidth
15 required by the consecutive freeze frames, which are inserted in between genuine I
16 frames and can typically be made very small in size in comparison to the average size of
17 a genuine I frame. Therefore, the insertion of 1, 2, 3 freeze frames will result in
18 bandwidth reductions of 2 to 1, 3 to 1 and 4 to 1 respectively. The accuracy of this
19 approximation degrades as more consecutive freeze frames and/or SNR scaling is
20 employed.) An easy way to see the validity of these approximate figures is to note for
21 example that in the case of 1 freeze frame insertion, the total presentation time of the
22 trick-mode clip for an originally 60 minutes duration asset will increase from 4 minutes
23 to 8 minutes. Since due to the underlying assumption of the first-order approximation

1 stated above, the same amount of data (I frames only) will be transmitted in this doubled
2 time interval, the bandwidth requirement will be halved. The final choice for trick-mode
3 implementation should reflect a balanced trade-off along these two degrees of freedom.
4 For example, SNR scaling of I frames down to 9 AC coefficients can be used along with
5 single freeze frame insertion between I frames. These two choices, both of which are
6 individually capable of providing a 2 to 1 bandwidth reduction as discussed before, will
7 yield a combined 4 to 1 bandwidth reduction which will comfortably bring the non-scaled
8 I frame-only bit-rate of 13,516,374.1 bits/sec. down to below the 4 Mbits/sec. quota. If
9 the visual quality provided by 9 AC coefficients is not considered adequate, then SNR
10 scaling could be tuned to keep more AC coefficients at the expense of a smaller
11 bandwidth reduction. This, however, could be compensated consequently by increasing
12 the number of freeze frames to be used in between I frames. Coarser quantization (and
13 therefore poorer visual quality) can be tolerated at high trick-mode speeds and better
14 visual quality should be retained at lower trick-mode speeds.

15 With reference to FIG. 2, if the client has requested trick-mode operation,
16 execution branches from step 58 to step 59. In step 59, execution branches to step 60 for
17 a low value of speed-up. In step 60, the trick-mode stream is produced by streaming
18 original-quality I frames and inserting three freeze frames per I frame, to yield a speed-up
19 factor of $15/4 = 3.75$ based on an original MPEG-2 coded stream having one I frame for
20 every 15 frames. For a higher speed-up factor, execution branches from step 59 to step
21 61. In step 61, either one or two freeze frames are selected per I frame to provide a
22 speed-up factor of $15/2 = 7.5$, or $15/3 = 5$ respectively. Then in step 62 the trick-mode
23 stream is produced by streaming reduced-quality I frames and inserting the selected

number of freeze frames between the reduced-quality I frames. If a trick-mode operation is not requested in step 58, then execution continues from step 58 to step 63. In step 63, the stream server computer streams original-quality MPEG-2 coded data to the client. Further details regarding trick-mode operation are described below with reference to FIGs. 7 to 10.

FIGs. 3 to 6 show further details regarding use of the present invention for MPEG splicing. In particular, reduced-quality frames are substituted for the freeze frames used in the seamless splicing procedure found in the common disclosure of Peter Bixby et al., U.S. application Ser. 09/539,747 filed March 31, 2000; Daniel Gardere et al., U.S. application Ser. 09/540,347 filed March 31, 2000; and John Forecast et al., U.S. application Ser. 09/540,306 filed March 31, 2000; which are all incorporated by reference herein. The common disclosure in these U.S. applications considered pertinent to the present invention is included in the written description below with reference to FIGs. 3 to 6 in the present application (which correspond to FIGs. 19, 22, 23, and 24 in each of the cited U.S. applications).

FIG. 3 shows a basic procedure for MPEG splicing. In the first step 121, the splicing procedure receives an indication of a desired end frame of the first clip and a desired start frame of the second clip. Next, in step 122, the splicing procedure finds the closest I frame preceding the desired start frame to be the In Point for splicing. In step 123, the splicing procedure adjusts content of the first clip near the end frame of the first clip and adjusts content of the second clip near the In Point in order to reduce presentation discontinuity (due to decoder buffer underflow) and also to prevent decoder buffer overflow when decoding the spliced MPEG stream. Finally, in step 124, the

1 concatenation of the first clip up to about the Out Point and the second clip subsequent to
 2 about the In Point is re-formatted, including re-stamping of the presentation time stamps
 3 (PTS), decoding time stamps (DTS), and program clock reference (PCR) values for the
 4 audio and video streams in the second clip.

5 Considering now video splicing, the splicing procedure should ensure the absence
 6 of objectionable video artifacts, preserve the duration of the spliced stream, and if
 7 possible, keep all of the desired frames in the spliced stream. The duration of the spliced
 8 stream should be preserved in order to prevent any time drift in the scheduled play-list.
 9 In some cases, it is not possible to keep all of the original video frames due to buffer
 10 problems.

11 Management of the video buffer is an important consideration in ensuring the
 12 absence of objectionable video artifacts. In a constant bit rate (CBR) and uniform picture
 13 quality sequence, subsequent pictures typically have coded representations of drastically
 14 different sizes. The encoder must manage the decoder's buffer within several constraints.
 15 The buffer should be assumed to have a certain size defined in the MPEG-2 standard.
 16 The decoder buffer should neither overflow nor underflow. Furthermore, the decoder
 17 cannot decode a picture before it receives it in full (i.e. completely). Moreover, the
 18 decoder should not be made to "wait" for the next picture to decode; this means that
 19 every 40 ms in PAL and 1/29.97 second in NTSC, the decoder must have access to a full
 20 picture ready to be decoded.

21 The MPEG encoder manages the video decoder buffer through decode time
 22 stamps (DTS), presentation time stamps (PTS), and program clock reference (PCR)
 23 values. When splicing the end of a first clip to the beginning of a second clip, there will

1 be a problem of video buffer management if a duration of time $DTS_{L1}-T_e$ is different from
2 a duration of time $DTS_{F2}-PCR_{e2}$ minus one video frame (presentation) interval, where
3 DTS_{L1} is the DTS at the end of the first clip and indicates the time at which the video
4 decoder buffer is emptied of video data from the first clip, T_e is the time at which the last
5 video frame's data is finished being loaded into the video decoder buffer, DTS_{F2} is the
6 DTS of the first frame of the second clip, and PCR_{e2} is the PCR of the second clip
7 extrapolated from the value of the most recent received genuine PCR record, to the first
8 byte of the picture header sync word of the first video frame in the clip to start. The
9 extrapolation adjusts this most recently received genuine PCR record value by the
10 quotient of the displacement in data bits of the clip from the position where it appears in
11 the second clip to the position at which video data of the first frame of the second clip
12 begins, divided by the data transmission bit rate for transmission of the clip to the
13 decoder. Because the time PCR_{e2} must immediately follow T_e , there will be a gap in the
14 decoding and presentation of video frames if $DTS_{F2}-PCR_{e2}$ is substantially greater than
15 $DTS_{L1}-T_e$ plus one video frame interval. In this case, the buffer will not be properly full
16 to begin decoding of the second clip one video frame interval after the last frame of the
17 first clip has been decoded. Consequently, either the second clip will be prematurely
18 started to be decoded or the decoder will be forced to repeat a frame one or more times
19 after the end of the display of the last frame from the first clip to provide the required
20 delay for the second clip's buffer build-up. In the case of a premature start for decoding
21 the second clip, a video buffer underflow risk is generated. On the other hand, in case of
22 repeated frames, the desired frame accuracy for scheduled play-lists is lost besides the
23 fact that neither a precise timing adjustment can be achieved through this procedure.

1 If $DTS_{F2}-PCR_{e2}$ is substantially less than $DTS_{L1}-T_e$ plus one video frame interval,
2 then the decoder will not be able to decode the first frame of the second clip at the
3 specified time DTS_{F2} because the last frame of the first clip will not yet have been
4 removed from the video buffer. In this case a video buffer overflow risk is generated.
5 Video buffer overflow may present a problem not only at the beginning of the second
6 clip, but also at a subsequent location of the second clip. If the second clip is encoded by
7 an MPEG-2 compliant encoder, then video buffer underflow or buffer overflow will not
8 occur at any time during the decoding of the clip. However, this guarantee is no longer
9 valid if the $DTS_{F2}-PCR_{e2}$ relationship at the beginning of the second clip is altered.
10 Consequently, to avoid buffer problems, the buffer occupancy at the end of the first clip
11 must be modified in some fashion. This problem is inevitable when splicing between
12 clips having significantly different ending and starting buffer levels. This is why the
13 Society of Motion Picture and Television Engineers (SMPTE) has defined some splice
14 types corresponding to well-defined buffer levels. (See SMPTE Standard 312M, entitled
15 "Splice Points for MPEG-2 Transport Streams," SMPTE Journal, Nov. 1998.) In order to
16 seamlessly splice the first clip to the second clip, the content of the first clip (towards its
17 end) is modified so that PCR_{e2} can immediately follow T_e (by one byte transmission
18 time) and DTS_{F2} can just follow DTS_{L1} (by one video frame presentation interval).

19 FIG. 4 shows a flow chart of a seamless video splicing procedure that attains the
20 desired condition just described above. In a first step 141, the first DTS of the second
21 clip is anchored at one frame interval later than the last DTS of the first clip in order to
22 prevent a video decoding discontinuity. Then, in step 142, the procedure branches
23 depending on whether the PCR extrapolated to the beginning frame of the second clip

1 falls just after the ending time of the first clip. If so, then the splice will be seamless with
2 respect to the original video content. Otherwise, the procedure branches to step 143. In
3 step 143, the content of the first clip is adjusted so that the PCR extrapolated to the
4 beginning frame of the second clip falls just after the ending time of the first clip.
5 Therefore the desired conditions for seamless video splicing are achieved.

6 With reference to FIG. 5, there is shown a more detailed flow chart of a seamless
7 video splicing procedure. In a first step 151, the procedure inspects the content of the
8 first clip to determine the last DTS/PTS of the first clip. This last DTS/PTS of the first
9 clip is designated DTS_{L1} . Next, in step 152, the procedure inspects the content of the first
10 clip to determine the time of arrival (T_e) of the last byte of the first clip. In step 153, the
11 procedure adds one frame interval to DTS_{L1} to find the desired first DTS location for the
12 second clip. The sum, designated DTS_{F1} , is equal to $DTS_{L1} + 1/FR$, where FR is the video
13 frame rate. In step 154, while keeping the DTS-PCR_e relationship unaltered for the
14 second clip, the procedure finds the time instant, designated T_s , at which the first byte of
15 the second clip should arrive at the decoder buffer. This is done by calculating
16 $T_{START} = DTS_{F2} - PCR_{e2}$, and $T_s = DTS_{F1} - T_{START}$.

17 Continuing in FIG. 6, in step 155, execution branches depending on whether T_s is
18 equal to T_e plus 8 divided by the bit rate. If not, then the clips to be spliced need
19 modification before concatenation, and execution branches to step 156. In step 156,
20 execution branches depending on whether T_s is less than T_e plus 8 divided by the bit rate.
21 If not, then there is an undesired gap in between the clips to be spliced, and execution
22 branches to step 157. In step 157, null packets are inserted into the clips to be spliced to
23 compensate for the gap. The gap to be compensated has a number of bytes, designated

1 G_r , equal to $(T_s - T_e)(\text{BIT RATE})/8$ minus one. If in step 156, T_s is less than T_e plus 8
2 divided by the bit rate, then execution continues from step 156 to step 158 to open up a
3 certain amount of space in the first clip to achieve $T_s = T_e + 8/(\text{BIT RATE})$. The number of
4 bytes to drop is one plus $(T_e - T_s)(\text{BIT RATE})/8$. If possible, the bytes are dropped by
5 removing null packets. Otherwise, one or more frames at the end of the first clip are
6 replaced with corresponding reduced-quality frames, which have fewer bytes than the
7 original-quality frames at the end of the first clip.

8 If in step 155 T_s is found to be equal to T_e plus 8 divided by the bit rate, then
9 execution continues to step 159. Execution also continues to step 159 from steps 157 and
10 158. In step 159, the transport streams from the two clips are concatenated. Finally, in
11 step 160, a subroutine is called to compute a video time stamp offset, designated as
12 V_{OFFSET} . This subroutine finds the DTS of the last video frame (in decode order) of the
13 first clip. This DTS of the last video frame of the first clip is denoted DTS_{VL1} . Then the
14 subroutine finds the original DTS of the first frame to be decoded in the second clip.
15 This DTS of the first frame to be decoded in the second clip is denoted DTS_{VF2} . Finally,
16 the subroutine computes the video time stamp offset V_{OFFSET} as $\text{DTS}_{\text{VL1}} - \text{DTS}_{\text{VF2}}$ plus one
17 video frame duration.

18 FIGs. 7 to 10 show further details regarding trick-mode operation. FIG. 7 shows
19 a timing relationship between video presentation units (VPUs) and associated audio
20 presentation units (APUs) in an original MPEG-2 coded data stream, and FIG. 8 shows
21 similar timing for the fast-forward trick-mode stream produced from the original data
22 stream of FIG. 7. (The fast-forward trick-mode stream is an example of a trick-mode
23 stream that could be produced in step 60 of FIG. 2.) The original data stream has

1 successive video presentation units for video frames of type I, B, B, P, B respectively.
 2 The trick-mode stream has successive video presentation units for video frames of types
 3 I, F, F, I, F where "F" denotes a freeze P (or possibly B) frame. Each I frame and
 4 immediately following F frames produce the same video presentation units as a
 5 respective I frame in the original data stream of FIG. 7, and in this example, one in every
 6 15 frames in the original data stream is an I frame. Each freeze frame is coded, for
 7 example, as a P frame repeating the previous I frame or the previous P-type freeze-frame
 8 (in display order). In each freeze frame, the frame is coded as a series of maximum-size
 9 slices of macroblocks, with an initial command in each slice indicating that the first
 10 macroblock is an exact copy of the corresponding macroblock in the previous frame
 11 (achieved by predictive encoding with a zero valued forward motion compensation vector
 12 and no encoded prediction error), and two consequent commands indicating that the
 13 following macroblocks in the slice until and including the last macroblock of the slice are
 14 all coded in the same way as the first macroblock.

15 For trick-mode operation, there is also a problem of how to select audio
 16 presentation units (APU) to accompany the video presentation units that are preserved in
 17 the trick-mode stream. Because the video presentation units (VPU) have a duration of
 18 $(1/29.97)$ sec. or about 33.37 msec. and the audio presentation units (APU) have a
 19 duration of 24 msec., there is neither a one-to-one correspondence nor alignment between
 20 VPUs and APUs. In a preferred implementation, the audio content of a trick-mode clip is
 21 constructed as follows. Given the total presentation duration $(1/29.97)$ sec. or about
 22 33.37 msec. for a single video frame, it is clear that always at least one and at most two
 23 24 msec. long audio presentation units (APU) will start being presented during the end-

1 to-end presentation interval of each video frame. This statement refers to the original clip
2 and does not consider any audio presentation unit whose presentation is possibly
3 continuing as the video frame under consideration is just put on the display. The first of
4 the above mentioned possibly two audio presentation units will be referred to as the
5 aligned audio presentation unit with respect to the video frame under consideration. For
6 example, in FIG. 8, the APU_j is the aligned audio presentation unit with respect to the
7 VPU_i . Now, when the I frames are extracted and possibly SNR scaled and possibly
8 further interleaved with a number of freeze P frames in between them to produce the
9 trick-mode video packetized elementary stream (PES), the associated trick-mode audio
10 stream is constructed as follows. For each I type video frame presentation interval (and
11 for that matter also for freeze P type video frames) in this trick-mode clip, the above
12 stated fact of at least one (and at most two) audio presentation unit being started, holds.
13 Then for each I frame presentation interval in the trick-mode clip, once any possibly
14 previously started and continuing audio presentation unit ends, insert its aligned audio
15 presentation unit (from the original clip) and continue inserting APUs from the original
16 clip subsequent to the aligned one until covering the rest of the I frame presentation
17 interval and also any possibly following freeze P frame presentation intervals until
18 crossing into and overlapping (or less likely aligning) with the next I frame's presentation
19 interval. In FIG. 8, for example, the audio presentation units APU_j , APU_{j+1} , APU_{j+2} , and
20 APU_{j+3} are inserted, until crossing into and overlapping with the next I frame VPU_{i+15} .
21 Following APU_{j+3} is inserted APU_k , which designates the APU aligned with VPU_{i+15} in
22 the original stream. Clearly, the final alignment of (the aligned and consequent) audio
23 presentation units with respect to their associated I frames will be slightly different in the

1 trick-mode clip as compared to the original clip. However, considering how the trick-
2 mode audio component will sound like, this poses no problem at all.

3 FIG. 9 is a flowchart of a procedure for producing the desired sequencing of audio
4 presentation units (APUs) in the fast-forward trick-mode stream. This procedure scans
5 the audio elementary stream in the original MPEG-2 stream to determine the sequence of
6 APUs in the original stream and their presentation-time alignment with the I frames in the
7 video elementary stream of the original MPEG-2 transport stream, while selecting APUs
8 to include in the trick-mode stream. In a first step 171, execution proceeds once the end
9 of the current APU is reached. If the end of the current APU has not entered a new VPU
10 (*i.e.*, the beginning of the current APU is within the presentation time of one VPU and the
11 end of the current APU is within the presentation time of the same VPU), or if it has
12 entered a new VPU (*i.e.*, the beginning of the current APU is within the presentation time
13 of one VPU and the end of the current APU is within the presentation time of a new
14 (next) VPU) but the new VPU is not an I frame, then execution branches to step 174. In
15 step 174, an APU pointer is incremented, and in step 175 execution proceeds into this
16 next APU. If in step 173 the end of the current APU extends into an I frame, then in step
17 176 the APU pointer is advanced to point to the first APU beginning within the duration
18 of the VPU of the I frame in the original MPEG-2 stream.

19 FIG. 10 is a flowchart of a procedure for producing a trick-mode stream from an
20 MPEG-2 transport stream (TS). In a first step 181, the MPEG-2 TS is inputted. In step
21 182, the video elementary stream (VES) is extracted from the TS. In step 183, a
22 concurrent task extracts the audio elementary stream (AES) from the TS. In step 184, I
23 frames are extracted from the VES and valid packetized elementary stream (PES) packets

1 are formed encapsulating the I frames. In step 185, the I frames are SNR scaled, for the
 2 high speed cases of the trick-mode. In step 186, P-type freeze frames are inserted into the
 3 stream of SNR scaled I frames (in between the scaled I frames), and valid PES packets
 4 are formed for the trick-mode VES encapsulating the P-type freeze frames and the SNR
 5 scaled I frames. Concurrently, in step 187, appropriate audio access units (from the
 6 originally input MPEG-2 TS asset) are selected and concatenated based on the structure
 7 of the VES being formed for the trick-mode clip, as described above with reference to
 8 FIG. 9, and valid PES packet encapsulation is formed around these audio access units.
 9 Finally, in step 188, the trick-mode TS stream is generated by multiplexing the trick-
 10 mode VES from step 186 into a system information (SI) and audio PES carrying TS
 11 skeleton including the audio PES packets from step 187.

12 FIGS 11 to 19 include details of the preferred techniques for truncating AC DCT
 13 coefficients for producing low-quality MPEG coded video from original-quality MPEG-2
 14 coded video. Most of these techniques exploit the fact that in the typical (default) zig-zag
 15 scan order, the basis functions for the high-order AC DCT coefficients have an increasing
 16 frequency content. FIG 11, for example, shows a matrix of the DCT coefficients C_{ij} . The
 17 row index (i) increases with increasing vertical spatial frequency in a corresponding 8x8
 18 coefficient block, and the column index (j) increases with increasing horizontal spatial
 19 frequency in the corresponding 8x8 coefficient block. The coefficient C_{11} has zero
 20 frequency associated with it in both vertical and horizontal directions, and therefore it is
 21 referred to as the DC coefficient of the block. The other coefficients have non-zero
 22 spatial frequencies associated with their respective basis functions, and therefore they are
 23 referred to as AC coefficients. Each coefficient has an associated basis function $f_{ij}(x,y)$

that is separable into x and y components such that $f_{ij}(x,y) = f_i(y)f_j(x)$. The x and y component functions $f_i(y)$ and $f_j(x)$ are shown graphically in FIG. 11 as cosine functions in order to illustrate their associated spatial frequencies. In practice, the component functions are evaluated at discrete points for the 64 pixel positions in the 8x8 blocks, so that each of the DCT basis functions is an 8x8 array of real numbers. In particular, the component functions are:

$$f_i(y) = \text{SQRT}((2 - \delta_{i-1})/8) (\cos((\pi/8)(y-1/2)(i-1))) \text{ for } y=1, 2, 3, \dots, 8$$

$$f_j(x) = \text{SQRT}((2 - \delta_{j-1})/8) (\cos((\pi/8)(x-1/2)(j-1))) \text{ for } x=1, 2, 3, \dots, 8$$

The heavy black line through the matrix of coefficients in FIG. 11 denotes the default zig-zag scan order typically used for MPEG-2 encoding. Listed in this order, the coefficients are $C_{11}, C_{12}, C_{21}, C_{31}, C_{22}, C_{13}, C_{14}, C_{23}, C_{32}, C_{41}, \dots, C_{86}, C_{77}, C_{68}, C_{78}, C_{87}, C_{88}$. The first coefficient in this zig-zag scan order is the DC coefficient C_{11} providing the lowest spatial frequency content in the 8x8 block of pixels, and the last coefficient in this zig-zag scan order is the coefficient C_{88} providing the highest spatial frequency content in the 8x8 block of pixels.

FIG. 12 is a diagram illustrating a relationship between an original MPEG-2 coded bit stream 200 and a reduced-quality MPEG-2 coded bit stream 210 resulting from truncation of high-order DCT coefficients from the original MPEG-2 coded bit stream. Shown in the original MPEG-2 coded bit stream 200 is a portion of a video PES packet including DCT coefficients for an 8x8 pixel block. The DCT coefficients include a differentially coded DC coefficient 201, and three (run, level) events 202, 203, 204

1 encoding three respective nonzero AC coefficients possibly along with some zero valued
2 AC coefficients preceding the three nonzero valued ones. The DCT coefficients are
3 ordered according to the zig-zag scan order shown in FIG. 11 (or possibly according to an
4 alternate zig-zag scan pattern also supported by the MPEG-2 standard), and AC
5 coefficients having zero magnitude are described in terms of total counts of consecutive
6 zero valued coefficients lying in between two nonzero valued coefficients, in the MPEG-
7 2 coded bit stream. An end-of-block (EOB) code 205 signals the end of the encoded
8 DCT coefficients for the current block. The reduced-quality MPEG-2 coded bit stream
9 210 includes a DC coefficient 201' identical to the DC coefficient 201 in the original
10 MPEG-2 coded bit stream 200, and a (run, level) event 202' identical to the (run, level)
11 event 202 in the original MPEG-2 coded bit stream 200. Second and third (run, level)
12 events, however, have been omitted from the reduced-quality MPEG-2 bit stream 210,
13 because an EOB code 205' immediately follows the (run, level) event 202'. Therefore,
14 the two nonzero high-order AC DCT coefficients encoded by the second and third (run,
15 level) events 203, 204 have been omitted from the reduced-quality MPEG-2 bit stream
16 210.

17 FIG. 13 is a flowchart of a procedure for scaling MPEG-2 coded video using a
18 variety of techniques including the omission of AC DCT coefficients. The procedure
19 operates upon an original-quality MPEG-2 coded video stream by removing AC DCT
20 coefficients in this stream to produce a lower quality MPEG coded video stream. In a
21 first step 221, execution branches to step 222 if the scaled MPEG coded video is to be
22 spatially subsampled. In step 222, the procedure removes any and all DCT coefficients
23 for spatial frequencies in excess of the Nyquist frequency for the downsampled video.

For example, if the low-quality video stream will be downsampled by a factor of two in both the vertical and the horizontal directions, then the procedure removes any and all DCT coefficients having a row index (i) greater than four and any and all DCT coefficients having a column index (j) greater than four. This requires the decoding of the (run, level) coded coefficients to the extent necessary to obtain an indication of the coefficient indices. If a sufficient number of the original AC DCT coefficients are removed for a desired bandwidth reduction, then the scaling procedure is finished. Otherwise, execution branches from step 223 to step 224. Execution also continues from step 221 to step 224 if spatial downsampling is not intended.

In step 224, execution branches to step 225 if low-pass scaling is desired. Low-pass scaling requires the least computational resources and may produce the best results if the scaled, low-quality MPEG coded video is spatially downsampled. In step 225, the procedure retains up to a certain number of lowest-order AC DCT coefficients for each block and removes any additional DCT coefficients for each block. This is a kind of frequency domain signal-to-noise ratio scaling (FDSNR) that will be designated FDSNR_LP. A specific example of the procedure for step 225 will be described below with reference to FIG. 14.

Execution continues from step 224 to step 226 if low-pass scaling is not desired. In step 226, execution branches to step 227 if largest magnitude based scaling is desired. Largest magnitude based scaling produces the least squared error or difference between the original-quality MPEG-2 coded video and the reduced-quality MPEG coded video for a given number of nonzero AC coefficients to preserve, but it requires more computational resources than the low-pass scaling of step 225. More computational

resources are needed because if there are more nonzero AC coefficients than the desired number of AC coefficients for a block, then the (run, level) events must be decoded fully to obtain the coefficient magnitudes, and additional resources are required to find the largest magnitude coefficients. In step 227, the procedure retains up to a certain number of largest magnitude AC DCT coefficients for each block, and removes any and all additional AC DCT coefficients for each block. This is a kind of frequency domain signal-to-noise ratio scaling (FDSNR) that will be designated FDSNR_LM. A specific example of the procedure for step 227 will be described below with reference to FIG. 15.

If in step 226 largest magnitude based scaling is not desired, then execution continues to step 228. In step 228, execution branches to step 229 to retain up to a certain number of AC DCT coefficients that differ in magnitude from up to that number of largest magnitude AC DCT coefficients by no more than a certain limit. This permits a kind of approximation to FDSNR_LM in which an approximate search is undertaken for the largest magnitude AC DCT coefficients if there are more nonzero AC DCT coefficients than the desired number of AC DCT coefficients in a block. The approximate search can be undertaken using a coefficient magnitude classification technique such as a hashing technique, and the low-pass scaling technique can be applied to the classification level that is incapable of discriminating between the desired number of largest magnitude AC DCT coefficients. A specific example is described below with reference to FIG. 19.

With reference to FIG. 14, there is shown a flowchart of a procedure for scaling MPEG-2 coded video using the low-pass frequency-domain signal-to-noise (FDSNR_LP) scaling technique. This procedure scans and selectively copies components of an input

stream of original-quality MPEG-2 coded data to produce an output stream of reduced-quality MPEG-2 coded video. The procedure is successively called, and each call processes coefficient data in the input stream for one 8x8 block of pixels. No more than a selected number “k” of coded lowest order (nonzero or zero valued) AC coefficients are copied for the block where the parameter “k” can be specified for each block.

In a first step 241 of FIG. 14, the procedure parses and copies the stream of original-quality MPEG-2 coded data up to and including the differential DC coefficient variable-length code (VLC). Next, in step 242, a counter variable “l” is set to zero. In step 243, the procedure parses the next (run, level) event VLC in the stream of original-quality MPEG-2 coded data. In step 244, if the VLC just parsed is an end-of-block (EOB) marker, execution branches to step 245 to copy the VLC to the stream of reduced-quality MPEG-2 coded video, and the procedure is finished for the current block.

In step 244, if the VLC just parsed is not an EOB marker, then execution continues to step 246. In step 246, a variable “r” is set equal to the run length of zeroes for the current (run, level) event, in order to compute a new counter value $l+r+1$. In step 247, if the new counter value $l+r+1$ is greater than the parameter “k”, then the procedure branches to step 248 to copy an EOB marker to the stream of reduced-quality MPEG coded data. After step 248, execution continues to step 249, where the procedure parses the input stream of original-quality MPEG-2 coded data until the end of the next EOB marker, and the procedure is finished for the current block.

In step 247, if the new counter value $l+r+1$ is not greater than the parameter “k”, then execution continues to step 250. In step 250, execution branches to step 251 if the new counter value $l+r+1$ is not equal to “k” (which would be the case if the new counter

1 value is less than “k”). In step 251, the counter state l is set equal to the new counter
2 value $l+r+1$. Then, in step 252, the VLC just parsed (which will be a VLC encoding a
3 (run, level) event) is copied from the stream of original-quality MPEG-2 coded data to
4 the stream of reduced-quality MPEG-2 coded data. After step 252, execution loops back
5 to step 243 to continue the scanning of the stream of original-quality MPEG-2 coded
6 data.

7 In step 250, if the new counter value $l+r+1$ is equal to “k”, then execution
8 branches from step 250 to step 253, to copy the VLC just parsed (which will be a VLC
9 encoding a (run, level) event) from the stream of original-quality MPEG-2 coded data to
10 the stream of reduced-quality MPEG-2 coded data. Next, in step 254, the procedure
11 copies an EOB marker to the stream of reduced-quality MPEG-2 coded data. After step
12 254, execution continues to step 249, where the procedure parses the input stream of
13 original-quality MPEG-2 coded data until the end of the next EOB marker, and the
14 procedure is finished for the current block.

15 FIG. 15 is a flowchart of a procedure for scaling MPEG-2 coded video using the
16 largest magnitude based frequency-domain signal-to-noise ratio (FDSNR_LM) scaling
17 technique. This routine is successively called, and each call processes coefficient data in
18 the input stream for one 8x8 block of pixels. No more than a specified number “k” of
19 largest magnitude AC DCT coefficients are copied for the block, and a different number
20 “k” can be specified for each block.

21 In a first step 261 in FIG. 15, the procedure parses and copies the input stream of
22 original-quality MPEG-2 coded data to the output stream of lower-quality MPEG-2 data
23 up to and including the differential DC coefficient variable-length code (VLC). Then in

1 step 262 all (run, level) event VLCs are parsed and decoded until and including the EOB
 2 marker of the current block. The decoding produces coefficient identifiers and
 3 corresponding quantization indices representing the quantized coefficient values. In step
 4 263, the quantization indices are transformed to quantized coefficient values. In step 264,
 5 the (quantized) coefficients are sorted in descending order of their magnitudes. In step
 6 265, the first “k” coefficients of the sorted list are preserved and the last 63-k AC DCT
 7 coefficients in the sorted list are set to zero. In step 266, (run, level) event formation and
 8 entropy coding (VLC encoding) are applied to the new set of coefficient values. Finally,
 9 in step 267, the VLCs resulting from step 266 are copied to the output stream until and
 10 including the EOB marker.

11 The sorting step 264 of the FDSNR_LM procedure can consume considerable
 12 computational resources. It is important to notice that not a full sorting of the quantized
 13 AC coefficients with respect to their magnitudes but rather a search for a specified
 14 number “k” of largest magnitude AC coefficients is all that is required. This task can be
 15 performed exactly or approximately in different ways so as to avoid the complexity
 16 associated with a conventional sorting procedure. In general, a relatively large number of
 17 the 63 AC DCT coefficients will have a quantized value of zero. Only the non-zero
 18 coefficients need be included in the sorting process. Moreover, if there are “n” non-zero
 19 coefficients and only “k” of them having the largest magnitudes are to be preserved in the
 20 output stream, then the sorting process may be terminated immediately after only the
 21 largest magnitude “k” coefficients have been found, or equivalently immediately after
 22 only the smallest magnitude “n-k” coefficients have been found. Moreover, the sorting

1 procedure itself can be different depending on a comparison of "k" to "n" in order to
2 minimize computations.

3 With reference to FIG. 16, there is shown a flowchart of a procedure that selects
4 one of a number of techniques for finding a certain number "k" of largest values out of a
5 set of "n" values. In a first step 271, execution branches to step 272 if "k" is less than $\frac{1}{2}$
6 "n." In step 272, execution branches to step 273 if "k" is much less than $\frac{1}{2}$ "n." In step
7 273, the first "k" values are sorted to produce a list of "k" sorted values, and then the last
8 "n-k" values are scanned for any value greater than the minimum of the sorted "k"
9 values. If a value greater than the minimum of the sorted "k" values is found, then that
10 minimum value is removed and the value greater than the minimum value is inserted into
11 the list of "k" sorted values. At the end of this procedure, the list of sorted "k" values
12 will contain the maximum "k" values out of the original "n" values. A specific example
13 of this procedure is described below with reference to FIG. 17.

14 In step 272, if "k" is not much less than $\frac{1}{2}$ "n", then execution branches to step
15 274. In step 274, a bubble-sort procedure is used, including "k" bottom-up bubble-sort
16 passes over the "n" values to put "k" maximum values on top of a sorting table. An
17 example of such a bubble-sort procedure is listed below:

18
19 /* TABLE(0) to TABLE(n-1) INCLUDES n VALUES */
20 /* MOVE THE k LARGEST OF THE n VALUES IN TABLE TO THE RANGE
21 TABLE(0) TO TABLE(k-1) IN THE TABLE */
22 /* k <= $\frac{1}{2}$ n */
23 FOR i=1 to k

```

1  FOR j=1 to n-i
2  IF (TABLE(n-j) > TABLE(n-j-1)) THEN(
3      /* SWAP TABLE(n-j) WITH TABLE(n-j-1) */
4      TEMP ← TABLE(n-j)
5      TABLE(n-j) ← TABLE(n-j-1)
6      TABLE(n-j-1) ← TEMP
7  NEXT j
8  NEXT I

```

In step 271, if “k” is not less than $\frac{1}{2}$ “n”, then execution branches to step 275. In step 275, if “k” is much greater than $\frac{1}{2}$ “n”, then execution branches to step 276. In step 276, a procedure similar to step 273 is used, except the “n-k” minimum values are maintained in a sorted list, instead of the “k” maximum values. In step 276, the last “n-k” values are placed in the sort list and sorted, and then the first “k” values are scanned for any value less than the maximum value in the sorted list. If a value less than the maximum value in the sorted list is found, then the maximum value in the sorted list is removed, and the value less than this maximum value is inserted into the sorted list. At the end of this procedure, the values in the sorted list are the “n-k” smallest values, and the “k” values excluded from the sorted list are the “k” largest values.

In step 275, if “k” is not much greater than $\frac{1}{2}$ “n”, then execution branches to step 277. In step 277, a bubble-sort procedure is used, including “n-k” top-down bubble-sort passes over the “n” values to put “n-k” minimum values at the bottom of a sorting table.

1 Consequently, the k maximum values will appear in the top “k” entries of the table. An
2 example of such a bubble-sort procedure is listed below:

3
4 /* TABLE(0) to TABLE(n-1) INCLUDES n VALUES */
5 /* MOVE THE n-k SMALLEST OF THE n VALUES IN THE TABLE */
6 /* TO THE RANGE TABLE(k) TO TABLE(n-1) IN THE TABLE */
7 /* $n > k \geq \frac{1}{2} n$ */
8 FOR i=1 to n-k
9 FOR j=0 to n-i-1
10 IF (TABLE(j) < TABLE(j+1)) THEN(
11 /* SWAP TABLE(j) WITH TABLE(j+1)*/
12 TEMP ←TABLE(j)
13 TABLE(j) ← TABLE(j+1)
14 TABLE(j+1) ← TEMP
15 NEXT j
16 NEXT I

17
18 Turning now to FIG. 17, there is shown a flowchart of a procedure for finding up
19 to a specified number “k” of largest magnitude AC DCT coefficients from a set of “n”
20 coefficients, corresponding to the procedure of FIG. 16 for the case of $k \ll \frac{1}{2}n$. In a first
21 step 281, a counter “i” is set to zero. In step 282, the next AC DCT coefficient is
22 obtained from the input stream of original-quality MPEG-2 coded data. If an EOB
23 marker is reached, as tested in step 283, then execution returns. In step 284, the counter

1 “i” is compared to the specified number “k”, and if “i” is less than “k”, execution
2 continues to step 285. In step 285, a coefficient index and magnitude for the AC DCT
3 coefficient is placed on a sort list. In step 286, the counter “i” is incremented, and
4 execution loops back to step 282.

5 Once the sort list has been loaded with indices and magnitudes for “k” AC DCT
6 coefficients and one additional coefficient has been obtained from the input stream,
7 execution branches from step 284 to step 287. In step 287 the list is sorted by magnitude,
8 so that the minimum magnitude appears at the end of the list. Then in step 288 the
9 coefficient magnitude of the current coefficient last obtained from the input stream is
10 compared to the magnitude at the end of the list. If the coefficient magnitude of the
11 current coefficient is not greater than the magnitude appearing at the end of the list, then
12 execution continues to step 289 to get the next AC DCT coefficient from the input
13 stream. If an EOB marker is reached, as tested in step 290, then execution returns.
14 Otherwise, execution loops back to step 288.

15 In step 288, if the magnitude of the current coefficient is greater than the
16 magnitude at the end of the list, then execution branches to step 291. In step 291, the
17 entry at the end of the list is removed. In step 292, a binary search is performed to
18 determine the rank position of the magnitude of the current coefficient, and in step 293,
19 the current coefficient index and magnitude are inserted into the list at the rank position.
20 The list, for example, is a linked list in the conventional fashion to facilitate the insertion
21 of an entry for the current coefficient at any position in the list. After step 293, execution
22 loops back to step 288.

An approximation technique of coefficient magnitude classification can be used to reduce the computational burden of sorting by coefficient magnitude. A specific example is the use of hashing of the coefficient magnitude and maintaining lists of the indices of coefficients having the same magnitude classifications. As shown in FIG. 18, a hash table 300 is linked to hash lists 301 storing the indices of classified coefficients. As shown, the hash table 300 is a list of 2^M entries, where "M" is three, and an entry has a value of zero if its associated list is empty, and otherwise the entry has a pointer to the end of the coefficients in its associated list. The lists shown in FIG. 18 have fixed memory allocations in which the pointers in the hash table also indicate the number of coefficient indices in the respective hash lists. Alternatively, the hash lists could be dynamically allocated and linked in the conventional fashion.

FIG. 19 shows a flowchart of a procedure for using the hash table 300 and hash lists 301 of FIG. 18 to perform a sort of "k" coefficients having approximately the largest magnitudes from a set of "n" coefficients. This approximation technique ensures that none of the "k" coefficients selected will have a magnitude that differs by more than a certain error limit from the smallest magnitude value of "k" coefficients having the largest magnitude. The error limit is established by the number of hash table entries, and it is the range of the magnitudes that can be hashed to the same hash table entry.

In a first step 311 in FIG. 19, the hash table is cleared. Then in step 312, the next AC DCT coefficient is obtained from the input stream. If an EOB marker is not reached, as tested in step 313, then execution continues to step 314. In step 314, a hash table index is stripped from the most significant bits (MSBs) of the coefficient magnitude. For the hash table in FIG. 18 having eight entries, the three most significant bits of the

coefficient magnitude are stripped from the coefficient magnitude. This is done by a bit masking operation together with a logical arithmetic shift operation. Then in step 315, the coefficient index is inserted on the hash list of the indexed hash table entry. For example, the hash table entry is indexed to find the pointer to where the coefficient index should be inserted, and then the pointer in the hash table entry is incremented. After step 315, execution loops back to step 312. Once all of the AC coefficients for the block have been classified by inserting them in the appropriate hash lists, an EOB marker will be reached, and execution will branch from step 313 to step 316.

Beginning in step 316, the hash table and hash lists are scanned to find approximately the “k” largest magnitude coefficients. The hash lists linked to the bottom entries of the hash table will have the indices for the largest magnitude coefficients. Each hash list is scanned from its first entry to its last entry, so that each hash list is accessed as a first-in-first-out queue. Therefore, in each magnitude classification, the coefficient ordering in the output stream will be the same as the coefficient ordering in the input stream, and the approximation will have a “low pass” effect in which possibly some lower-frequency coefficients having slightly smaller magnitudes will be retained at the expense of discarding some higher-frequency coefficients having slightly larger magnitudes. (The approximation results from the fact that the last hash list to be scanned is not itself sorted, and to eliminate the error of the approximation, the last hash list to be scanned could be sorted.)

In step 316, a scan index “i” is set to 2^M-1 in order to index the hash table beginning at the bottom of the table, and a counter “j” is set equal to “k” in order to stop the scanning process after finding “k” coefficients. Next, in step 317, the hash table is

1 indexed with “i”. In step 318, if the indexed entry of the hash table is zero, then
2 execution branches to step 319. In step 319, the procedure is finished if “i” is equal to
3 zero; otherwise, execution continues to step 320. In step 320, the index “i” is
4 decremented, and execution loops back to step 317.

5 If in step 318 the indexed hash table entry is not zero, then execution continues to
6 step 321. In step 321, the next entry is obtained from the indexed hash list, and the
7 coefficient index in the entry is used to put the indexed coefficient in the output stream.
8 Then in step 322 execution branches to step 319 if the end of the indexed hash list is
9 reached in the previous step 321. If the end of the list was not reached in step 321, then
10 execution continues from step 322 to step 323. In step 323 the counter “j” is
11 decremented, and in step 324 the counter “j” is compared to zero. In step 324, if the
12 counter “j” is less than or equal to zero, then the procedure is finished. Otherwise, if the
13 counter “j” is not less than or equal to zero in step 324, execution loops back to step 321.

14 The FDSNR_LM procedure, as described above, in general provides a significant
15 improvement in peak signal-to-noise ratio (PSNR) over the FDSNR_LP procedure when
16 each procedure retains the same number of non-zero AC DCT coefficients. It has been
17 found, however, that substantially more bits are required for the (run, level) coding of the
18 non-zero AC DCT coefficients resulting from the FDSNR_LM procedure than those
19 resulting from the FDSNR_LP procedure, provided that the same coefficient quantization
20 and scanning method is used. Therefore, the FDSNR_LM procedure provides at best a
21 marginal improvement in rate-distortion (PSNR as a function of bit rate) over the
22 FDSNR_LP procedure unless the non-zero AC DCT coefficients for the FDSNR_LM
23 procedure are quantized, scanned, and/or (run, level) coded in a fashion different from the

1 quantization, scanning, and/or (run, level) coding of the coefficients in the original
2 MPEG-2 clip. A study of this problem resulted in a discovery that it is sometimes
3 possible to reduce the number of bits for (run, level) coding of coefficients for an 8x8
4 block including a given number of the non-zero largest magnitude AC DCT coefficients
5 if additional coefficients are also (run, level) coded for the block.

6 The (run, level) coding of the non-zero AC DCT coefficients from the
7 FDSNR_LM procedure has been found to require more bits than from the FDSNR_LP
8 procedure due to an increased occurrence frequency of escape sequences for the (run,
9 level) coding. The increased frequency of escape sequences is an indication that the
10 statistical likelihood of possible (run, level) combinations for the non-zero AC DCT
11 coefficients selected by the FDSNR_LM procedure is different from the statistical
12 likelihood of possible (run, level) combinations for the non-zero AC DCT coefficients
13 produced by the standard MPEG-2 coding process and in particular those selected by the
14 FDSNR_LP procedure.

15 The MPEG-2 coding scheme assigns special symbols to the (run, level)
16 combinations that occur very frequently in ordinary MPEG-2 coded video. The most
17 frequent (run, level) combinations occur for short run lengths (within the range of about
18 0 to 5, where the run length can range from 0 to 63) and relatively low levels (about 1 to
19 10, where the level can range from 1 to 2048). The most frequent of these special
20 symbols are assigned the shortest variable-length code words (VLCs). If a (run, level)
21 combination does not have such a special symbol, then it is coded as an escape sequence
22 including a 6-bit escape sequence header code word followed by a 6-bit run length
23 followed by a 12 bit signed level. An escape sequence requires a much greater number of

bits than the special symbols, which have varying lengths depending on their relative frequency. In particular, the escape sequences each have 24 bits, and the special symbols have variable-length code words having a maximum of 17 bits.

There are two (run, level) VLC tables. The first coding table is designated TABLE 0, and the second is designated TABLE 1. These tables specify the (run, level) combinations having special symbols, and the special symbol for each such combination. For each table, the (run, level) combinations having special symbols, and the range of the VLC bit lengths of the special symbols, are summarized below:

SUMMARY OF PROPERTIES OF DCT COEFFICIENT TABLE ZERO

(Table Zero is Table B.14, p. 135 of ISO/IEC 13818-2 1996E)

Run	Range of Levels	Range of Code Lengths
0	1 to 40	2 to 16
1	1 to 18	4 to 17
2	1 to 5	5 to 14
3	1 to 4	6 to 14
4	1 to 3	6 to 13
5	1 to 3	7 to 14
6	1 to 3	7 to 17
7	1 to 2	7 to 13
8	1 to 2	8 to 13
9	1 to 2	8 to 14

1	10	1 to 2	9 to 14
2	11	1 to 2	9 to 17
3	12	1 to 2	9 to 17
4	13	1 to 2	9 to 17
5	14	1 to 2	11 to 17
6	15	1 to 2	11 to 17
7	16	1 to 2	11 to 17
8	17	1	13
9	18	1	13
10	19	1	13
11	20	1	13
12	21	1	13
13	22	1	14
14	23	1	14
15	24	1	14
16	25	1	14
17	26	1	14
18	27	1	17
19	28	1	17
20	29	1	17
21	30	1	17
22	31	1	17

23

1 SUMMARY OF PROPERTIES OF DCT COEFFICIENT TABLE ONE

2 (Table One is Table B.15, p. 139 of ISO/IEC 13818-2 1996E)

3

4 Run Range of Levels Range of Code Lengths

5	0	1 to 40	3 to 16
6	1	1 to 18	4 to 17
7	2	1 to 5	6 to 14
8	3	1 to 4	6 to 14
9	4	1 to 3	7 to 13
10	5	1 to 3	7 to 14
11	6	1 to 3	8 to 17
12	7	1 to 2	8 to 13
13	8	1 to 2	8 to 13
14	9	1 to 2	8 to 14
15	10	1 to 2	8 to 14
16	11	1 to 2	9 to 17
17	12	1 to 2	9 to 17
18	13	1 to 2	9 to 17
19	14	1 to 2	10 to 17
20	15	1 to 2	10 to 17
21	16	1 to 2	11 to 17
22	17	1	13
23	18	1	13

1	19	1	13
2	20	1	13
3	21	1	13
4	22	1	14
5	23	1	14
6	24	1	14
7	25	1	14
8	26	1	14
9	27	1	17
10	28	1	17
11	29	1	17
12	30	1	17
13	31	1	17

14

15 The FDSNR_LP procedure selected AC DCT coefficients have (run, level)
16 symbol statistics that are similar to the statistics of ordinary MPEG-2 coded video, and
17 therefore the FDSNR_LP AC DCT coefficients have a similar frequency of occurrence
18 for escape sequences in comparison to the ordinary MPEG-2 coded video. In contrast,
19 the FDSNR_LM procedure selects AC DCT coefficients resulting in (run, level)
20 combinations that are less likely than the combinations for ordinary MPEG-2 coded
21 video. This is due to two reasons. First, the FDSNR_LM procedure selects AC DCT
22 coefficients having the highest levels. Second, the FDSNR_LM procedure introduces
23 higher run lengths due to the elimination of coefficients over the entire range of

coefficient indices. The result is a significantly increased rate of occurrence for escape sequences. Escape sequences form the most inefficient mode of coefficient information encoding in MPEG-2 incorporated into the standard so as to cover important but very rarely occurring coefficient information.

In order to improve the rate-distortion performance of the scaled-quality MPEG-2 coded video from the FDSNR_LM procedure, the non-zero AC DCT coefficients selected by the FDSNR_LM procedure should be quantized, scanned, and/or (run, level) coded in such a way that tends to reduce the frequency of the escape sequences. For example, if the original-quality MPEG-2 coded video was (run, level) coded using TABLE 0, then the largest magnitude coefficients should be re-coded using TABLE 1 because TABLE 1 provides shorter length VLCs for some (run, level) combinations having higher run lengths and higher levels. It is also possible that re-coding using the alternate scan method instead of the zig-zag scan method may result in a lower frequency of occurrence for escape sequences. For example, each picture could be (run, level) coded for both zig-zag scanning and alternate scanning, and the scanning method providing the fewest escape sequences, or the least number of bits total, could be selected for the coding of the reduced-quality coded MPEG video.

There are two methods having general applicability for reducing the frequency of escape sequences resulting from the FDSNR_LM procedure. The first method is to introduce a non-zero, "non-qualifying" AC DCT coefficient of the 8x8 block into the list of non-zero qualifying AC DCT coefficients to be coded for the block. In this context, a "qualifying" coefficient is one of the k largest magnitude coefficients selected by the FDSNR_LM procedure. The non-qualifying coefficient referred to above, must be lying

1 in between two qualifying AC DCT coefficients (in the coefficient scanning order) that
 2 generate the (run, level) combination causing the escape sequence. Moreover, this non-
 3 qualifying coefficient must cause the escape sequence to be replaced with two shorter
 4 length VLCs when the AC DCT coefficients are (run, level) coded. This first method has
 5 the effect of not only decreasing the number of bits in the coded reduced-quality MPEG
 6 video in most cases, but also increasing the PSNR.

7 The qualifying AC DCT coefficient causing the escape sequence that is first in the
 8 coefficient scanning order will be simply referred to as the first qualifying coefficient.
 9 The qualifying AC DCT coefficient causing the escape sequence that is second in the
 10 coefficient scanning order will be simply referred to as the second qualifying coefficient.
 11 For example, suppose the qualifying coefficients in zig-zag scan order for an 8x8 block
 12 include C_{51} followed by C_{15} having a level of 40. If only the qualifying coefficients were
 13 (run, level) coded for the microblock, C_{15} would result in a run length of 3, because there
 14 are a total of three non-qualifying coefficients (C_{42} , C_{33} , and C_{24}) between C_{51} and C_{15} in
 15 the scan order. Therefore, C_{15} would have to be coded as an escape sequence, because a
 16 run of 3 and level of 40 does not have a special symbol. In this example, the escape
 17 sequence is in effect caused by a first qualifying coefficient, which is C_{51} , and a second
 18 qualifying coefficient, which is C_{15} . This escape sequence can possibly be eliminated
 19 say, if C_{24} is a non-zero, non-qualifying coefficient of the block, C_{24} has a level of 5 or
 20 less, and C_{24} is (run, level) coded together with the qualifying coefficients. For example,
 21 assuming that C_{24} has a level of 5, and using the MPEG-2 (run, level) coding TABLE 1,
 22 then C_{24} has a run length of two and is coded as the special symbol 0000 0000 1010 0s,
 23 where "s" is a sign bit, and C_{15} now has a run length of 0 and is coded as the special

1 symbol 0000 0000 0010 00s. Such a consideration clearly applies to the rest of the non-
2 zero non-qualifying coefficients lying in between the two qualifying coefficients
3 producing the escape sequence. In the above example, these non-qualifying coefficients
4 are C_{42} and C_{33} .

5 Whether or not an escape sequence can be eliminated from the (run, level) coding
6 of the qualifying coefficients can be determined by testing a sequence of conditions. The
7 first condition is that the second qualifying coefficient must have a level that is not
8 greater than the maximum level of 40 for the special (run, level) symbols. If this
9 condition is satisfied, then there must be a non-zero non-qualifying AC DCT coefficient
10 that is between the first and second qualifying coefficients in the coefficient scanning
11 order. If there is such a non-qualifying coefficient, then the combination of its level and
12 the run length between the first qualifying coefficient and the non-qualifying coefficient
13 in the coefficient scanning order must be one of the special (run, level) symbols. If so,
14 then the combination of the level of the second qualifying coefficient and the run length
15 between the non-qualifying coefficient and the second qualifying coefficient must also be
16 a special (run, level) symbol, and if so, all required conditions have been satisfied. If not,
17 then the conditions with respect to the non-qualifying coefficient are successively applied
18 to any other non-zero non-qualifying AC DCT coefficient of the block lying in between
19 the two qualifying coefficients, until either all conditions are found to be satisfied or all
20 such non-qualifying coefficients are tested and failed. If there are sufficient
21 computational resources, this search procedure should be continued to find all such non-
22 qualifying coefficients that would eliminate the escape sequence, and to select the non-

1 qualifying coefficient that converts the escape sequence to the pair of special symbols
2 having respective code words that in combination have the shortest length.

3 A flow chart for a modified FDSNR_LM procedure using the first method is
4 shown in FIGS. 20 and 21. In a first step 331 of FIG. 20, the procedure finds up to "k"
5 largest magnitude non-zero AC DCT coefficients (i.e., the "qualifying coefficients") for
6 the block. (This first step 331 is similar to steps 261 to 265 of FIG. 15, as described
7 above.) In step 332, (run, level) coding of the qualifying coefficients is begun in the scan
8 order using the second coding table (Table 1). This (run, level) coding continues until an
9 escape sequence is reached in step 333, or the end of the block is reached in step 336. If
10 an escape sequence is reached, execution branches from step 333 to step 334. If the level
11 of the second qualifying coefficient causing the escape sequence is greater than 40,
12 execution continues from step 334 to step 336. Otherwise, execution branches from step
13 334 to step 335 to invoke a subroutine (as further described below with reference to FIG.
14 21) to possibly include a non-zero non-qualifying AC DCT coefficient in the (run, level)
15 coding to eliminate the escape sequence. The subroutine either returns without success,
16 or returns such a non-qualifying coefficient so that the escape sequence is replaced with
17 the two new (run, level) codings of the first qualifying coefficient and the non-qualifying
18 coefficient and then the non-qualifying coefficient and the second qualifying coefficient.
19 From step 335, execution continues to step 336. Execution returns from step 336 if the
20 end of the block is reached. Otherwise, execution continues from step 336 to step 337, to
21 continue (run, level) coding of the qualifying coefficients in the scan order using the
22 second coding table (TABLE 1). This (run, level) coding continues until an escape

sequence results, as tested in step 333, or until the end of the block is reached, as tested in step 336.

With reference to FIG. 21, there is shown a flow chart of the subroutine (that was called in step 335 of FIG. 20) for attempting to find a non-zero, non-qualifying AC DCT coefficient that can be (run, level) coded to eliminate an escape sequence for a qualifying coefficient. In a first step 341, the procedure identifies the first qualifying coefficient and the second qualifying coefficient causing the escape sequence. For example, the subroutine of FIG. 21 can be programmed as a function having, as parameters, a pointer to a list of the non-zero AC DCT coefficients in the scan order, an index to the first qualifying coefficient in the list, and an index to the second qualifying coefficient in the list. In step 342, the subroutine looks for a non-zero non-qualifying AC DCT coefficient between the first and the second qualifying coefficients in the scan order. For example, the value of the index to the first qualifying coefficient is incremented and compared to the value of the index for the second qualifying coefficient, and if they are the same, there is no such non-qualifying coefficient. Otherwise, if the new coefficient pointed to (by incrementing the index of the first qualifying coefficient) is a non-zero coefficient then it becomes a candidate non-qualifying coefficient deserving further testing. If however the new coefficient pointed to (by incrementing the index of the first qualifying coefficient) has a value zero then it is not a candidate non-qualifying coefficient. If no such (candidate) non-qualifying coefficients are found, as tested in step 343, then execution returns from the subroutine with a return code indicating that the search has been unsuccessful. Otherwise, execution continues to step 344.

1 In step 344, the non-qualifying coefficient is (run, level) coded, to determine in
 2 step 345 whether it codes to an escape sequence. If it codes to an escape sequence, then
 3 execution loops back from step 345 to step 342 to look for another non-zero non-
 4 qualifying AC DCT coefficient in the scan order between the first and second qualifying
 5 coefficients. If it does not code to an escape sequence, then execution continues from
 6 step 345 to step 346. In step 346, the second qualifying coefficient is (run, level) coded,
 7 using the new run length, which is the number of coefficients in the scan order between
 8 the non-qualifying coefficient and the second qualifying coefficient. If it codes to an
 9 escape sequence, as tested in step 347, then execution loops back from step 347 to step
 10 342 to look for another non-zero non-qualifying AC DCT coefficient in the scan order
 11 between the first and second qualifying coefficients. If it does not code to an escape
 12 sequence, then execution continues from step 347 to step 348.

13 In step 348, execution returns with a successful search result unless a continue
 14 search option has been selected. If the continue search option has been selected, then
 15 execution branches from step 348 to step 349 to search for additional non-zero non-
 16 qualifying AC DCT coefficients that would eliminate the escape sequence. In other
 17 words, steps 342 to 347 are repeated in an attempt to find additional non-zero non-
 18 qualifying AC DCT coefficients that would eliminate the escape sequence. If no more
 19 such non-qualifying coefficients are found, as tested in step 350, execution returns with a
 20 successful search result. Otherwise, execution branches from step 350 to step 351 to
 21 select the non-qualifying coefficient giving the shortest overall code word length and/or
 22 the largest magnitude for the best PSNR, and execution returns with a successful search
 23 result. For example, for each non-qualifying coefficient that would eliminate the escape

sequence, the total bit length is computed for the (run, level) coding of the non-qualifying coefficient and the second qualifying coefficient. Then a search is made for the non-qualifying coefficient producing the shortest total bit length, and if two non-qualifying coefficients which produce the same total bit length are found, then the one having the largest level is selected for the elimination of the escape sequence.

A second method of reducing the frequency of occurrence of the escape sequences in the (run, level) coding of largest magnitude AC DCT coefficients for an 8x8 block is to change the mapping of coefficient magnitudes to the levels so as to reduce the levels. Reduction of the levels increases the likelihood that the (run, level) combinations will have special symbols and therefore will not generate escape sequences. This second method has the potential of achieving a greater reduction in bit rate than the first method, because each escape sequence can now be replaced by the codeword for one special symbol, rather than by the two codewords as is the case for the first method. The second method, however, may reduce the PSNR due to increased quantization noise resulting from the process producing the lower levels. Therefore, if a desired reduction of escape sequences can be achieved using the first method, then there is no need to perform the second method, which is likely to reduce the PSNR. If the first method is used but not all of the escape sequences have been eliminated, then the second method could be used to possibly eliminate the remaining escape sequences.

The mapping of coefficient magnitudes to the levels can be changed by decoding the levels to coefficient magnitudes, changing the quantization scale factor (qsi), and then re-coding the levels in accordance with the new quantization scale factor (qsi). The quantization scale factor is initialized in each slice header and can also be updated in the

1 macroblock header on a macroblock basis. Therefore it is a constant for all blocks in the
2 same macroblock. In particular, the quantization scale factor is a function of a
3 q_scale_type parameter and a quantizer_scale_code parameter. If q_scale_type = 0, then
4 the quantizer scale factor (qsi) is twice the value of q_scale_code. If q_scale_type =1,
5 then the quantizer scale factor (qsi) is given by the following table, which is the right half
6 of Table 7-6 on page 70 of ISO/IEC 13838-2:1996(E):

8	<u>quantizer scale code</u>	<u>quantization scale factor (qsi)</u>
9	1	1
10	2	2
11	3	3
12	4	4
13	5	5
14	6	6
15	7	7
16	8	8
17	9	10
18	10	12
19	11	14
20	12	16
21	13	18
22	14	20
23	15	22

1	16	24
2	17	28
3	18	32
4	19	36
5	20	40
6	21	44
7	22	48
8	23	52
9	24	56
10	25	64
11	26	72
12	27	80
13	28	88
14	29	96
15	30	104
16	31	112
17		

18 In a preferred implementation, to reduce the coefficient levels, the quantization
19 scale factor is increased by a factor of two, and the levels of the non-zero AC DCT
20 coefficients are reduced by a factor of two, so long as the original value of the
21 quantization scale factor is less than or equal to one-half of the maximum possible
22 quantization scale factor. For `q_scale_type = 1`, a factor of two increase in the

1 quantization scale factor (qsi) is most easily performed by a table lookup of a new
 2 quantization_scale_code using the following conversion table:

3

4	<u>Original quantization scale code</u>	<u>New quantization scale code</u>
5	1	2
6	2	4
7	3	6
8	4	8
9	5	9
10	6	10
11	7	11
12	8	12
13	9	14
14	10	16
15	11	17
16	12	18
17	13	19
18	14	20
19	15	21
20	16	22
21	17	24
22	18	25
23	19	26

1	20	27
2	21	28
3	22	29
4	23	30
5	24	31

6

7 In a preferred method for generation of trick mode files, the quantization scale
8 factor is adjusted in order to achieve a desired reduction in the escape sequence
9 occurrence frequency resulting from the modified FDSNR_LM procedure, and the
10 number (k) of largest magnitude coefficients is adjusted in order to achieve a desired
11 reduction in bit rate. A specific implementation is shown in the flow chart of FIGS. 22-
12 23. In a first step 361, the number (k) of largest magnitude AC coefficients per 8x8 block
13 is initially set to a value of 9, and the quantization scaling factor (QSF) is initially set to a
14 value of 2. Then conversion of the I frames of an original-quality MPEG-2 coded video
15 clip to a lower quality level begins. When a picture header is encountered in step 362,
16 indicating the beginning of a new I frame, execution continues to step 363. In step 363,
17 execution branches depending on the value of the intra_vlc_format parameter in the
18 picture header of the original-quality MPEG-2 coded video clip. This value is either 0,
19 indicating that the first (run, level) coding table (TABLE 0) was used for coding the
20 picture, or 1, indicating that the second (run, level) coding table (TABLE 1) was used for
21 coding the picture. In either case, the down scaled quality picture will be coded with the
22 second (run, level) coding table. If the intra_vlc_format parameter is equal to 0 execution
23 continues from step 363 to step 364 where TABLE 0 is read in for (run, level) symbol

1 decoding in the original-quality MPEG-2 coded clip. Otherwise, if the intra_vlc_format
2 parameter is equal to 1, then execution continues from step 363 to step 365 where
3 TABLE 1 is read in for (run, level) symbol decoding in the original-quality MPEG-2
4 coded clip.

5 After steps 364 and 365, execution continues to step 366. In step 366, the
6 modified FDSNRS_LM procedure is applied to the 8x8 blocks of the current slice, using
7 the adjusted quantization scale index, if the adjusted quantization scale index is less than
8 the maximum possible quantization scale index. In step 367, execution loops back to step
9 362 to continue 8x8 block conversion until a new slice header is encountered, indicating
10 the beginning of a new slice. Once a new slice is encountered, execution continues from
11 step 367 to step 368. In step 368, the average escape sequence occurrence frequency per
12 block for the last slice is compared to a threshold TH1. If the escape sequence
13 occurrence frequency is greater than the threshold, then execution branches to step 369.
14 In step 369, if the quantization scaling factor (QSF) is less than or equal to a limit value
15 such as 2, then execution branches to step 370 to increase the quantization scaling factor
16 (QSF) by a factor of two.

17 In step 368, if the escape sequence occurrence frequency is not greater than the
18 threshold TH1, then execution continues to step 371 of FIG. 23. In step 371, the average
19 escape sequence occurrence frequency per 8x8 block for the last slice is compared to a
20 threshold TH2. If the escape sequence occurrence frequency is less than the threshold
21 TH2, then execution branches to step 372. In step 372, if the quantization scaling factor
22 (QSF) is greater than or equal to a limit value such as 2, then execution branches to step
23 373 to decrease the quantization scaling factor (QSF) by a factor of two. After step 373,

1 and also after step 370 of FIG. 22, execution continues to step 374 of FIG. 23. In step
 2 374, execution continues to step 375 if a backtrack option has been selected. In step 375,
 3 re-coding for the last slice is attempted using the adjusted quantization scale factor. The
 4 new coding, or the coding that gives the best results in terms of the desired reduction of
 5 escape sequence occurrence frequency, is selected for use in the scaled quality picture.
 6 After step 375, execution continues to step 376. Execution also continues to step 376
 7 from: step 369 in FIG 22 if the quantization scaling factor (QSF) is not less than or equal
 8 to 2; step 371 in FIG 23 if the escape sequence occurrence frequency is not less than the
 9 threshold TH2; step 372 in FIG 23 if the quantization scaling factor (QSF) is not greater
 10 than or equal to 2; and from step 374 in FIG 23 if the backtrack option has not been
 11 selected.

12 In step 376, the average bit rate of the (run, level) coding per 8x8 block for at
 13 least the last slice is compared to a high threshold TH3. Preferably this average bit rate is
 14 a running average over the already processed segment of the current scaled quality I-
 15 frame, and the high threshold TH3 is selected to prevent video buffer overflow in
 16 accordance with the MPEG-2 Video Buffer Verifier restrictions. If the average bit rate
 17 exceeds the high threshold TH3, then execution continues to step 377, where the number
 18 (k) of non-zero largest magnitude AC coefficients per 8x8 block is compared to a lower
 19 limit value such as 6. If the number (k) is greater than or equal to 6, then execution
 20 continues to step 378 to decrement the number (k).

21 In step 376, if the average bit rate is not greater than the threshold TH3, then
 22 execution continues to step 379. In step 379, the average bit rate is compared to a lower
 23 threshold TH4. If the average bit rate is less than the threshold TH4, then execution

1 branches from step 379 to step 380, where the number (k) of non-zero largest magnitude
2 AC DCT coefficients per 8x8 block is compared to a limit value of 13. If the number (k)
3 is less than or equal to 13, then execution continues to step 381 to increment the number
4 (k). After step 378 or 381, execution continues to step 382. In step 382, execution
5 continues to step 383 if a backtrack option is selected. In step 383, an attempt is made to
6 re-code the last slice for the scaled quality picture using the adjusted value of the number
7 (k) of non-zero largest magnitude AC DCT coefficients per block. After step 383,
8 execution loops back to step 362 of FIG. 22 to continue generation of the scaled quality
9 clip. Execution also loops back to step 362 of FIG. 22 after: step 377 if the value of (k) is
10 not greater than or equal to 6; step 379 if the average bit rate is not less than the threshold
11 TH4; step 380 if the value of (k) is not less than or equal to 13; and step 382 if the
12 backtrack option has not been selected. Coding of the scaled quality clip continues until
13 the end of the original quality clip is reached in step 364 of FIG. 22, in which case
14 execution returns.

15 In a preferred implementation, a fast forward trick mode file and a fast reverse
16 trick mode file are produced from an original-quality MPEG-2 coded video main file
17 when the main file is ingested into the video file server. As shown in FIG. 24, a volume
18 generally designated 390 is allocated to store the main file 391. The volume 390 includes
19 an allocated amount of storage that exceeds the real file size of the main file 391 in order
20 to provide additional storage for meta-data 392, the fast forward trick file 393, and the
21 fast reverse trick file 394. The trick files are not directly accessible to clients as files;
22 instead, the clients may access them thorough trick-mode video service functions. With

1 this strategy, the impact on the asset management is a minimum. No modification is
2 needed for delete or rename functions.

3 Because the volume allocation is done once for the main file and its fast forward
4 and fast reverse trick mode files, there is no risk of lack of disk space for production of
5 the trick files. The amount of disk blocks to allocate for these files is computed by the
6 video service using a volume parameter (vsparams) specifying the percentage of size to
7 allocate for trick files. A new encoding type is created in addition to types RAW for
8 direct access and MPEG2 for access to the main file. The new encoding type is called
9 EMPEG2, for extended MPEG2, for reference to the main file plus the trick files. The
10 video service allocates the extra file size only for these files.

11 For the transfer of these files to archive or to another video file server, it would be
12 useful to transfer all the data even if it is a non-standard format. For the FTP copy-in, a
13 new option is added to specify if the source is in the EMPEG2 format or if it is a standard
14 MPEG2 file. In the first case, the copy-in should provide the complete file 390. In the
15 second case, the video service allocates the extra size and the processing is the same as
16 for a record. For the copy-out, the same option can be used to export the complete file
17 390 or only the main part 391. The archiving is always done on the complete file 390.

18 The trick mode file production is done by a new video service procedure. This
19 procedure takes as input the speed-up factor (or the target trick mode file size) along with
20 the number of freeze (P or B) frames to insert in between the scaled I frames and then
21 generates both the fast forward file 393 and the fast reverse file 394 for this speed-up
22 factor (or target trick mode file size) and with the specified number of interleaving freeze
23 frames. Since the bandwidth of the original clip (in the main file) and the bandwidths of

1 the two trick mode clips (in the fast forward and fast reverse files) are the same, the
2 speed-up factor and the target trick mode file size are equivalent pieces of information. A
3 default speed-up factor (system parameter) can be used. The main file is read and the
4 trick mode files are produced. If a trick mode file already exists with the same speed-up
5 factor, it is rewritten or nothing is done depending on an option. Multiple trick mode
6 files could be created with different speed-up factors. But it is preferred to permit only
7 one set of fast forward and fast reverse trick mode files to be produced at a time (i.e., no
8 parallel generation with different speed-up factors). The current speed-up factor is a
9 parameter of the volume parameters (vsparams).

10 As stated above another parameter to be provided to the video service procedure
11 in charge of trick mode file generation is the number of freeze frames to be inserted in
12 between consequent scaled I frames. The preferred values for this parameter are 0 and 1,
13 although other positive integer values greater than 1 are also possible. The inclusion of
14 freeze frames due to their very small sizes spare some bandwidth which can then be used
15 to improve the quality of scaled I frames. Hence, the freeze frames in this context provide
16 a mechanism to achieve a trade-off between the scaled I frame quality and the temporal
17 (motion) sampling. Depending on the speed-up factor (or the target trick mode file size)
18 and also the number of interleaving freeze frames to be inserted, the video service
19 procedure in charge of trick mode file generation determines a sub-sampling pattern
20 (closest to uniform) to choose the original I frames which will be scaled and included in
21 the trick mode files. For example, the case of an original clip with 10 frames per GOP, a
22 trick mode file size which is 10% of the main file together with 0 freeze frames, implies
23 the use of all original I frames for being scaled and included in the trick mode file. This

will typically result in a low quality scaling. As another example, the case of an original clip with 10 frames per GOP, a trick mode file size which is 10% of the main file together with 1 freeze frame, implies the use of a 2 to 1 (2:1) sub-sampling on the original I frames which will choose every other original I frame for being scaled and included in the trick mode file.

FIG. 25 is a more detailed diagram of the volume 390, showing additional meta-data and related data structures. The Inode 401 includes 4 disk blocks containing a file-system oriented description of the file. The Meta-data (MD) directory 402 includes 4 disk blocks describing each entry of the meta-data area 392. The entries of the meta-data area 392 include a description of the MPEG-2 meta-data 403, a description of the trick files header meta-data 404, and a description of the GOP index meta-data 405. The MPEG-2 meta-data 403 includes 15 disk blocks maximum.

The trick files header 404 includes 1 disk block, which specifies the beginning of free area (end of last trick file) in blocks, the number of trick files couple (FF FR), and for each trick file, a speed-up factor, a block address of the GOP index, a block address of the trick file forward, a byte length of the trick file forward, a block address of the trick file reverse, a byte length of the trick file reverse, a frames number of the trick file, and a number of GOP of each trick files.

The GOP index includes 2024 disk blocks. The GOP index specifies, for each GOP, a frame number, a pointer to the MPEG-2 data for the GOP in the main file, and various flags and other attributes of the GOP. The flags indicate whether the GOP entry is valid and whether the GOP is open or closed. The other attributes of the GOP include the maximum bit rate, the average bit rate, the AAU size in bytes, the APU duration in

1 seconds, the audio PES packet starting locations, the AAU starting locations, the AAU
 2 PTS values, and the decode time stamp (DTS) and the value of the program clock
 3 reference (PCR) extrapolated to the first frame of the GOP. The size of all the data
 4 preceding the main file is, for example, 1 megabyte.

5 There is one GOP index 406 for both the fast forward file 393 and the fast reverse
 6 file 394. The GOP index 406 of the trick files is different than the GOP index 405 of the
 7 main file. The GOP index 406 of the trick files contains, for each GOP, the byte offset in
 8 the trick file forward of the TS packet containing the first byte of the SEQ header, the
 9 frame number in the fast forward file of the GOP (the same value for the fast reverse file
 10 can be computed from this value for the fast forward file), the frame number in the
 11 original file of the first frame of the GOP, and the byte offset in the original file of the
 12 same frame (to resume after fast forward or reverse without reading the main GOP
 13 index).

14 The GOP index 405 for the main file and the GOP index 406 for the fast forward
 15 and fast reverse trick files provides a means for rapidly switching between the normal
 16 video-on-demand play operation during the reading of the main file, and the fast-forward
 17 play during the reading of the fast-forward file, and the fast-reverse play during the
 18 reading of the fast reverse file. For example, FIG. 26A illustrates the read access to
 19 various GOPs in the main file, fast forward file, and fast reverse file, during a play
 20 sequence listed in FIG. 26B. Due to the presence of down-scaled I frames and possibly
 21 consequent freeze frames in the trick mode files, the video buffer verifier (VBV) model
 22 for a trick mode file is different than the VBV model of the main file. Consequently, the
 23 mean video decoder main buffer fullness levels can be significantly different for these

It is desired to copy in and out of the volume 390 with or without the meta-data 392 and the trick files 393, 394. This is useful to export and/or import complete files without regenerating the trick files. The file encoding type is now recognized as a part of the volume name. Therefore there can be multiple kinds of access to these files. The read and write operations are done by derivations of the class file system input/output (FSIO) which takes into account the proper block offset of the data to read or write. There is one derivation of FSIO per encoding type, providing three different access modes. EMGP3, MPEG2, and RAW. EMPEG2 accesses the whole volume from the beginning of the meta-data array, and in fact provides access to the entire volume except the inode 401, but no processing is done. MPEG2 access only the main part of the asset

1 with MPEG processing, including file analyze and meta-data generation in a write access.
2 RAW access only the main part of the asset without processing. These access modes are
3 operative for read and write operations for various access functions as further shown in
4 FIG. 27.

5 During a record operation, the video service allocates a volume and computes the
6 number of block to allocate using the volume parameter giving the percentage to add for
7 the trick files. Then, the size in blocks given to the stream server is the main part size
8 only without the extension for the trick files. This avoids using the reserved part of the
9 volume when the effective bit rate is higher than the requested bit rate. At the end of a
10 record operation or a FTP copyin operation, the video service calls a procedure
11 CMSPROC_GETATTR, and the stream server returns the actual number of bytes
12 received and the actual number of blocks used by the main file plus the meta-data. The
13 same values are returned for both MPEG2 and EMPEG2 files. The video service
14 computes again the file extension to manage the trick files and adjust the number of
15 allocated blocks.

16 Both trick files forward and reverse are generated by the same command. First,
17 the trick file forward is generated by reading the main file. The trick file GOP index is
18 built and kept in memory. During this generation, only the video packets are kept. PCR,
19 PAT and PMT will be regenerated by the MUX in play as for any other streams. The
20 audio packets are discarded. This ensures that there is enough stuffing packets for the
21 PCR reinsertion. For this, a stuffing packet is inserted every 30 milliseconds.

22 Then using the GOP index, the trick file forward is read GOP by GOP in reverse
23 order to generate the trick file reverse. The same GOPs are present in both files. The

only modification done is an update of the video PTS, which must be continuous. Then, the GOP index is written on disk. This avoids reading again the file while generating the second trick file. The GOP index size is: 24 times the GOP number. In the worst case (the file is assumed not to be 1 frame only), there are 2 frames per GOP and 30 frames per second. So for 1 hour in fast forward, the GOP index size is: $(24 \times 3600 \times 30) / 2 = 1296000$ bytes. This will be the case for a 4 hour film played at 4 times the normal speed. Therefore, this GOP index can be kept in memory during the trick file generations without risk of memory overflow.

The read and write rate are controlled to conserve bandwidth on the cached disk array. The bandwidth reserved for these generations is a parameter given by the video service. It is a global bandwidth for both read and writes. The number of disk I/O per seconds is counted so as not to exceed this bandwidth.

13 The trick files header update is done once when both the fast forward and fast
14 reverse trick files and the GOP index have been successfully written.

15 Playing a file is done with the CM_MpegPlayStream class. Fast forward
16 (reverse) can only be requested when we are in the paused state. The current frame on
17 which we are paused is known from the MpegPause class. This frame is located in the
18 GOP index of the trick file. Then the clip start point and length are modified in the Clip
19 instance with the trick file position computed from the beginning of the clip. So, the Clip
20 class handle these trick files in a manner similar to the main file. The current logical
21 block number is updated with the block address in the trick file recomputed from the
22 beginning of the main clip. In fact, a seek is performed in the trick file as it was part of
23 the main file, which is totally transparent for the ClipList and Clip classes. The transition

from fast forward to pause is handled in a similar fashion. The clip start and length and the logical block number are again updated. The smooth transitions from pause to fast forward and from fast forward to pause are done in the same way as for regular play. There is a splicing from the pause stream to the play stream.

The class hierarchy for trick file handling is shown in FIG. 28. The MpegFast, MpegFastForward and MpegFastReverse class handles the GOP generation from the initial file. This is the common procedure for building the GOP whatever the source and the destination. RealTimeFastFwd and RealTimeFastRev are the class instantiated when a real time fast forward (reverse) has to be done. They manage the real-time buffer flow to the player. There is a derivation of the methods takeBuffer and returnBuffer which uses the base class to build the GOP in the buffer to be played. The main file access is done using a buffer pool.

TrickFilesGenerate is the class instantiated to generate trick files forward and reverse. It inherits from TrickFileAccess the methods for reading the original file some buffers and for writing the trick file and its meta-data. It inherits from MpegFastForward the methods for building the GOP and for managing the advance in the file.

The computation of the next 1 frame to play is done by MpegFast, MpegFastForward and RealTimeFastFwd. When a trick file generation command is invoked, a thread is created and started and the generation itself is done off-line. A call-back is sent to the video service when the generation is completed. The class TrickFilesGenerate generates the trick file forward, and then, using the GOP index built in memory, the class TrickFiles Generate generates the trick file reverse.

When there is a transition from play to pause, the only latency issue is related to the buffer queue handled by the player and to the GOP size. The stream can build immediately the active pause GOP, and then this GOP will be sent at the end of the current GOP with a splicing between these two streams.

When there are transitions from pause to regular play or fast forward and fast reverse, a seek in the file is done. This means that the current buffer pool content is invalidated and the buffer pool is filled again. Play can start again while the buffer pool is not completely full, as soon as the first buffer is read. The buffer pool prefilling can continue as a background process. The issue here is that there is a risk to generate an extra load on the cached disk array as well as on the stream server side when the buffer pool is being prefilled.

To avoid too frequent transitions from play to fast forward and fast reverse, there is a limitation of the number of requests per second for each stream. This limitation is part of the management of the video access commands. A minimum delay between two commands is defined as a parameter. If the delay between a request and the previous one is too small, the request is delayed. If a new request is received during this delay, the new request replaces the waiting one. So the last received request is always executed.

The volume parameter (vsparams) file contains these new parameters for the trick mode files:

TrickFileExtensionSize:<percent>:

DefaultFastAcceleration:<acceleration>:

DMtrickFileGen:<mask of reserved DM> (This parameter is a mask of the stream
servers that can be chosen to perform the trick file generation. The default value is
0xffff: all of the stream servers.)

DMtrickFileGenBW:<bandwidth used for trick file generation> (This parameter
is the value of the bandwidth effectively used by the stream server for the trick files
generation.)

The video service routines are modified to operate upon the EMPG2 files, and in particular to compute the size of the EMPG2 files, to allocate the volume for the main file and the trick files, and to generate the trick files. The volume creation functions (VAPP) and volume access functions (RRP) use the EMPG2 files in the same way as MPEG2 files. This means that a MPEG2 volume is created on the stream server. Both MPEG2 and EMPG2 files can be used in the same session or play-list. The session encoding type is MPEG2. In record (or copy-in), the number of blocks allocated for an EMPG2 file is computed using the percentage of size to add. At the end of record (or copy-in), the number of blocks is adjusted using the number of blocks returned by the stream server (by CMSPROC_GETATTR) and adding the percentage for trick files. The trick files validity and generation date are stored by the video service in the asset structure. The bandwidth allocated to the TrickFilesGenerate command is defined in the volume parameters (vsparams or vssiteparams). The selection of a stream server to generate the trick files takes into account this bandwidth only. If preferred stream servers are specified in vsparams (or vssiteparams), then the selected stream server will be one of these specified stream servers.

1 In a preferred implementation of the video service software, a new encoding type
2 is created. The encoding type enum becomes:

```
3       enum encoding_t{  
4           ENC_UNKNOWN   = 0,           /* unknown format */  
5           ENC_RAW       = 1,           /* uninterpreted data */  
6           ENC_MPEG1     = 2,           /* constrained MPEG1 */  
7           ENC_MPEG      = 3,           /* generic MPEG */  
8           ENC_EMPEG2    = 4,           /* MPEG2 with trick files extension */  
9       };
```

10
11 - The encoding information accessible by VCMP_EXTENDEDINFO includes
12 information about trick files:

```
13  
14   struct trickFilesInfo_t{  
15       ulong_t       generationDate;     /* date/time of the generation of the trick  
16                   files */  
17       rate_factor_t acceleration;       /* acceleration factor */  
18       ulong_t       framesNumber;       /* frames number in each trick file (FWD and  
19                   REV) */  
20       ulong_t       gopNumber;          /* GOP number of each file */  
21   };
```

```
22  
23   struct EMPEG2info_t{
```

```

1      MPEG2info_t      MPEG2info;
2      trickFilesInfo_t  trickFiles<>;
3  };
4
5  union encodingInfo_t switch (encoding-t enc){
6      case ENC_MPEG:
7          MPEG2info_t      MPEG2info;
8      case ENC_EMPEG2:
9          EMPEG2info_t      EMPEG2info;
10     default:
11         void;
12 };
13

```

The video service software includes a new procedure (VCMP_TRICKFILESGEN) for trick file generation, which uses the following structures:

```

17 struct VCMPtrickgenres_t{
18     VCMPstatus_t      status;
19     tHandle_t         handle;
20 };
21
22 struct VCMPtrickfilesargs_t{
23     name_t            clipname;

```

```

1      bool_t      overwriteIfExists;
2      rate_factor_t  acceleration;
3  };
4
5  VCMPtrickgenres_t      VCMP_TRICKFILESGEN (VCMPtrickfilesargs_t)  =
6  36,
7

```

8 If the trick files already exist and if the boolean overwriteIfExists is true, then the
 9 trick files are generated again, in the other case nothing is done. Acceleration is the
 10 acceleration as defined and used for the controlled speed play function. It is a percentage
 11 of the normal speed, it must be greater than 200 and smaller than 2000. The special value
 12 of 0 can be used to generate files with the default acceleration defined in vssiteparams.
 13 The procedure starts the generation process. The completion is notified by a callback.

14 The video service includes a new option to copyin and copyout. The option is
 15 added to allow a user to copy all the file or the main asset only. For compatibility with
 16 old client programs, the following new procedures are added:

```

17
18  VCMPcopyres_t      VCMP_FULL_COPYIN      (copyinargs2_t)      = 37,
19  VCMPcopyres_t      VCMP_FULL_COPYOUT     (copyoutargs2_t)     = 38,
20

```

21 These new procedures take the same interface as the existing one, but are used to copy-in
 22 the complete file: meta-data + Asset + trick files.

23

```

1      The video service includes a new procedure
2      VCMP_TRICKFILESGENCOMPLETED, which uses the following structures:
3
4      struct VCMPtrickfilescomplete_t{
5          tHandle_t          handle;
6          VCMPstatus_t      status;
7      };
8
9      VCMPstatus_t TRICKFILESGENCOMPLETED (VCMPtrickfilescomplete_t)    = 10,
10
11      The video service includes new procedures are added for handling trick mode
12      generation arguments, which uses the following structures:
13
14      struct cms_trick_gen_args {
15          Handle_t          Vshandle;
16          name_t            name;
17          bool_t            overwriteIfExists;
18          rate_factor_t     acceleration;
19          bandwidth_t       reservedBw;
20      };
21
22      cms_status    CMSPROC_GEN_TRICK_FILES (cms_trick_gen_args)          = 34,
23

```

```

1 struct trick_gen_completed_args {
2     Handle_t      Vshandle;
3     cms_status     status;
4 };
5 void CTLPROC_TRICKGENCOMPLETED (trick_gen_completed_args) = 8,
6

```

7 The video service includes the following option to force the regeneration of trick
8 files even if they exist:

```
9     nms_content -gentrick <name> [<-f>] [acceleration]
```

10 Without this option, an error code is returned if the trick files exist. "Acceleration" is an
11 acceleration factor. If it is not present, the default value is taken in vsparams.

12 The video services include a encoding information access function (nms_content
13 -m). This function produces a displayed output containing, for each trick file generated,
14 the acceleration, the generation date and time, the frames number, and the GOP number.

15 For the use of an FTP copy function with the trick files, the following new
16 commands are added:

```
17
18 nms_content -copyinfull <same arguments as -copyin>
```

```
19 nms_content -copyoutfull <same arguments as -copyout>
```

20
21 Another application of the SNR scaling of the invention is to reduce the bit rate of
22 an MPEG-2 transport stream in order to allow combining multiple MPEG-2 transport
23 streams to match a target bit rate for a multiple program transport stream. For example,

FIG. 29 shows a system for combining an MPEG-2 audio-visual transport stream 411 with an MPEG-2 closed-captioning transport stream 412 to produce a multiplexed MPEG-2 transport stream 413. In this case, the closed captioning transport stream 412, containing alphanumeric characters and some control data instead of audio-visual information, has a very low bit rate compared to the audio-visual transport stream 411. Assuming that the target bit rate for the multiplexed transport stream 413 is the same as the bit rate of the audio-visual transport stream 411, there need be only a slight decrease in the bit rate of the audio-visual transport stream, and this slight decrease can be obtained by occasionally removing one non-zero AC DCT coefficient per 8x8 block. Therefore, in the system of FIG. 29, the audio-visual transport stream 411 is processed by a program module 414 for selective elimination of non-zero AC DCT coefficients to slightly reduce the average bit rate of this transport stream. A transport stream multiplexer 415 then combines the modified audio-visual transport stream with the closed captioning transport stream 412 to produce the multiplexed MPEG-2 transport stream 413.

In order to determine whether or not any non-zero AC DCT coefficient should be eliminated from a next 8x8 block in the audio-visual transport stream 411, a module 421 is executed periodically to compute a desired bit rate change in the audio-visual transport stream 411. For example, respective bit rate monitors 416, 417 may measure the actual bit rate of the audio-visual transport stream 411 and the closed captioning transport stream 412. Alternatively, if it is known precisely how these transport streams are generated, presumed values for the bit rates of these transport streams may be used in lieu of measured bit rates. The computation of the desired bit rate change also includes the

desired bit rate 418 for the multiplexed MPEG-2 transport stream, and a bit rate 419 of multiplexer overhead, representing any net increase in bit rate related to the multiplexing of the audio-visual transport stream 411 with the closed captioning transport stream 412. An adder/subtractor 420 combines the various bit rate values from the inputs 416, 417, 418, and 419 to compute the desired bit rate change in the audio-visual transport stream 411. From the adder/subtractor 420, the module 421 converts the desired change in bit rate to a desired number of bits to be removed per computational cycle (e.g., per millisecond). This number of bits to be removed per computational cycle is received in an adder/subtractor 422, and the output of the adder/subtractor is received in an integrator 423. A limiter 424 takes the sign (positive or negative) of the integrated value to produce a flag indicating whether or not one non-zero AC DCT coefficient should be removed from the coefficients for the next 8x8 block, assuming that the next block has at least one non-zero AC DCT coefficient. (Alternatively, a non-zero AC DCT coefficient could be removed only if the 8x8 block has more than a predetermined fraction of the average number of AC DCT coefficients per 8x8 block.) The particular non-zero AC DCT coefficient to remove in each case can be selected using any of the methods discussed above with reference to FIGS. 14, 15, or FIG. 20. For example, the coefficient to remove could be the last non-zero AC DCT coefficient in the scan order. Alternatively, the non-zero AC DCT coefficient having the smallest magnitude could be removed so long as its removal does not cause an escape sequence.

When the module 414 removes a non-zero AC DCT coefficient from a 8x8 block, it sends the number of bits removed to the adder/subtractor 422. In a preferred implementation, the operations of the adder/subtractor 422, integrator 423, and limiter

1 424 are performed by a subroutine having a variable representing the integrated value.
2 During each computational cycle, the variable is incremented by the number of bits to be
3 removed per computational interval, and whenever the module 414 removes a non-zero
4 AC DCT coefficient from a 8x8 block of the audio-visual transport stream, the variable is
5 decremented by the number of bits removed.

6 Although the system in FIG. 29 has been described for achieving a slight
7 reduction in bit rate of the MPEG-2 audio-visual transport stream 411 for combining
8 multiple transport stream to produce a multiplexed MPEG-2 transport stream, it should
9 be apparent that it could be used for obtaining relatively large reductions in bit rate. In
10 this case, the module 414 would use the procedure of FIGS. 14, 15 or preferably FIG. 20,
11 and a multi-level comparator 424 would be used instead of a single-level comparator 424.
12 The multi-level comparator would determine a desired number of non-zero coefficients to
13 discard per 8x8 block based on the value of the output of the integrator 423. The
14 maximum number of non-zero AC coefficients to keep for each 8x8 block (i.e., the value
15 of the parameter "k"), for example, would be determined by subtracting the number of
16 non-zero AC DCT coefficients in the 8x8 block from the desired number to discard, and
17 limiting this difference to no less than a predetermined fraction of the average number of
18 non-zero AC coefficients per 8x8 block.

19 In order to produce reduced-quality MPEG-2 data having a bit rate that is a small
20 fraction of the bit rate of the original-quality MPEG-2 data, the bits for each frame should
21 be allocated to each DCT coded block (i.e., an 8x8 pixel block) in the frame in such a
22 way that more complex blocks are allocated more bits than simple blocks. With a
23 relatively large reduction in bit rate, there will be degradation in picture quality, and the

1 degradation will be less noticeable if each 8x8 block in the frame suffers approximately
2 the same degree of degradation. If the same number of bits is allocated to each 8x8 block
3 regardless of complexity, then the complex blocks will be allocated fewer bits than
4 needed for a uniform degradation, and the simple blocks will be allocated more bits than
5 needed for a uniform degradation. Consequently, the picture degradation will become
6 more noticeable in the more complex blocks.

7 In the original-quality MPEG-2 data, the number of bits allocated to the AC DCT
8 coefficients of each 8x8 block is a good estimate of the complexity of the block, and
9 therefore one can allocate coefficient bits to each 8x8 block in the reduced quality
10 MPEG-2 data in proportion to block complexity by proportionally scaling the number of
11 bits allocated to coefficient bits in each 8x8 block of the original quality MPEG-2 block.
12 For example, a coefficient bit rate reduction factor of RF is computed, where RF is the
13 ratio of the desired AC DCT coefficient bit rate of the reduced-quality MPEG-2 data
14 divided by the AC DCT coefficient bit rate of the original quality MPEG-2 data. The
15 number of coefficient bits available for encoding the non-zero AC DCT coefficients in
16 each 8x8 block in the reduced-quality MPEG-2 data is then calculated by multiplying the
17 rate reduction factor RF by the number of coefficient bits in the corresponding 8x8 block
18 in the original-quality MPEG-2 data.

19 The final reduced video bit rate is smaller than the total bit rate of the original-
20 quality MPEG-2 data, divided by the reduction factor, by an offset S representing the rate
21 of bits in the MPEG-2 data that are not video bits or padding bits. The padding bits may
22 be removed in the reduced rate file. The preferred rate reduction technique of dropping
23 DCT coefficients has an advantage that the reduced-quality MPEG-2 data will have the

1 same offset S of bits in the reduced rate MPEG-2 data that are not coefficient bits.
2 Therefore, if the original-quality MPEG-2 data has a total bit rate of BO , a padding PD
3 and offset S , and it is desired for the reduced-quality MPEG-2 data to have a total bit rate
4 of BR , then the coefficient bit reduction factor RF is $RF = (BR-S)/(BO-PD-S)$. $BR-S$
5 represents the video bit rate of the reduced bit rate file and $BO-PD-S$ represents the video
6 bit rate of the original quality MPEG-2 data. In many cases the offset S will be relatively
7 small compared to BR and BO so that any change or uncertainty in the value of the offset
8 has a small effect on the value of the reduction factor. Therefore, the offset S could be
9 estimated based on frame size and frame type (e.g., I, P or B) depending on the desired
10 uncertainty in the value of the reduction factor RF . S is always greater than zero.

11 In general, the number of bits allocated to the non-zero AC DCT coefficients of
12 an 8×8 block in a frame in the reduced-quality MPEG-2 data will be less than the number
13 of bits allocated to the non-zero AC DCT coefficients of the corresponding 8×8 block of
14 the same frame in the original-quality MPEG-2 data. This is due to the fact that the
15 preferred process of bit rate reduction is eliminating non-zero AC DCT coefficients.
16 Therefore, the number of bits allocated for the 8×8 block is reduced by a variable number
17 each time a coefficient is eliminated. It is possible although not likely that the dropping
18 of one or more coefficients will cause the number of bits allocated to the 8×8 block to
19 become equal to the number of bits that are available for encoding the non-zero AC DCT
20 coefficients for the reduced 8×8 block. It is most likely that some of the bits available for
21 encoding the non-zero AC DCT coefficients for the reduced 8×8 block will not be used
22 for encoding the non-zero AC DCT coefficients for the reduced 8×8 block. So that the
23 actual bit rate for the reduced-quality MPEG-2 data will not become significantly less

than the desired rate due to available bits that would not be used for encoding the non-zero AC DCT coefficients, the available bits that are not used for encoding non-zero AC DCT coefficients for a reduced 8x8 block are made available for encoding AC DCT coefficients in following 8x8 blocks. In particular, a count of the bits that were available but not yet used for encoding AC DCT coefficients for the reduced blocks is maintained in a memory location or program variable that will be referred to as a "bucket." Initially, this bucket is empty. At the start of selecting a reduced number of coefficients for the next reduced-quality 8x8 block, the coefficient bit rate reduction factor RF is multiplied by the number of coefficient bits in the original-quality MPEG-2 block. This product represents a minimum number of bits that are available for encoding non-zero AC DCT coefficients selected for the reduced-quality block. In addition, a certain fraction of any bits in the bucket are also available for encoding non-zero AC DCT coefficients selected for the reduce-quality block. The fraction represents the extent to which the bits in the bucket can be shared among following blocks. For example, if the bucket is not cleared and the fraction is the reciprocal of the number of blocks in a frame, then bits in the bucket will be shared over at least a frame of following coefficients. Moreover, at the end of a frame, the bucket need not be cleared, and excess bits from coefficients in one frame will be available for encoding AC DCT coefficients in a following frame.

FIG. 30 illustrates graphically how the adaptive bit rate method achieves a generally proportional reduction in the number of non-zero AC coefficient bits in each 8x8 block of a frame, and a smoothing of deviations from a proportional reduction among the 8x8 blocks. A frame 501 in the original-quality MPEG coded video is represented as a matrix of 8x8 blocks. In this example, for the sake of illustration, the frame has 8 rows

1 and 8 columns of 8x8 blocks, although in practice, the frame would usually have a larger
 2 number of rows and columns of 8x8 blocks. Inside the cells of the matrix of the frame
 3 501 are compressed 8x8 blocks, such as the compressed block 502. The compressed
 4 blocks are shown having a varying size. The size of each compressed block represents
 5 the total number of bits used in the variable-length encoding of the non-zero AC DCT
 6 coefficients in the block.

7 Shown below the frame 501 is a corresponding reduced size frame 503
 8 representing a frame in the reduced-quality MPEG coded video produced from the
 9 original-quality MPEG coded video. In particular, the reduced size frame 503 is a matrix
 10 having the same number of rows and columns of 8x8 blocks. Each 8x8 block in the
 11 reduced frame 503 is produced from a corresponding 8x8 block in the original frame 501
 12 by a process of selective removal of non-zero AC DCT coefficients so that the 8x8 block
 13 in the reduced frame 504 has a proportionately reduced number of bits encoding non-zero
 14 AC DCT coefficients, subject to smoothing of deviations among the reduced 8x8 blocks.
 15 For example, the reduced 8x8 block 504 is produced by selective removal of non-zero
 16 AC DCT coefficients from the compressed block 502.

17 It is desired to keep the ratio between each original block size and the reduced
 18 rate block size constant and equal to the ratio of the original frame size and the reduced
 19 bit rate frame size. This is not always possible because the bit allocation per block differ
 20 from frame to frame due to local simple blocks and a variable proportion of complex sub
 21 frames in a frame. For example, there are blocks for which the maximum number if bits
 22 is smaller than the average block size computed as:

23

1 $B_{reduced}(j) = B_{original}(j) / MA_{original}(n) * BR_{reduced}$

2

3 Where (j) is the index of the block in the frame, (n) is the present frame number,

4 $B_{reduced}(j)$ is the number of bits allocated for the reduced bit rate block (j) and it is
5 proportional to the original block size $B_{original}(j)$, $MA_{reduced}(n)$ is the moving average
6 of original frame size after (n) frames and $BR_{reduced}$ is the target reduced video bit rate.

7 The average block size of the reduced bit rate frame is calculated as:

8

9 $AVB_{reduced} = F_{reduced} / NB$

10

11 Where $F_{reduced}$ is the frame size of the reduced-quality file, and NB is the number of
12 blocks per frame and it is calculated to conform to the MPEG-2 standard from the frame
13 resolution as follows:

14

15 $NB = \text{ceil}(FR(\text{horizontal})/8) * \text{ceil}(FR(\text{vertical})/8)$

16

17 where "ceil" is the conversion to a rounded-up integer, $FR(\text{horizontal})$ is the horizontal
18 frame resolution in pixels (720 for NTSC and 528 for PAL), and $FR(\text{vertical})$ is the
19 vertical frame resolution in pixels (480 for NTSC and PAL).

20 In other cases the number of bits is greater than the average size allocated as
21 shown in the reduced frame 503. It should also be apparent from the above discussion
22 that sometimes all of the bits allocated to the reduced frame cannot be used. Bits that are

1 allocated but not used in the reduced frame should be distributed by a smoothing
2 mechanism to another reduced frame that may be more complex.

3 Fig 30 shows how some reduced blocks are smaller than a proportionally-reduced
4 size and the bucket is filled, and some reduced blocks are larger than the proportionally-
5 reduced size and the bucket is emptied. This methodology is especially efficient in cases
6 when there are very large variations between adjoining blocks based on the details of the
7 block, for example hair strings near a talking head. For example, the reduced frame 503
8 includes a “bucket filling” block 505 and a “bucket emptying” block 507. A rectangle
9 506 around the bucket filling block 505 indicates a proportional reduction in the number
10 of bits of the non-zero AC DCT coefficients in the corresponding compressed block 509
11 in the original frame 501. The difference in area between the outer rectangle 506 and the
12 bucket filling block 505 indicate bits that are available for encoding non-zero AC DCT
13 coefficients of the bucket filling block 505 but were not used and instead were placed in
14 the bucket. A rectangle 508 inside the bucket emptying block 507 indicates a
15 proportional reduction in the number of bits of the non-zero AC DCT coefficients in the
16 corresponding compressed block 510 in the original frame 501. The difference in area
17 between the inner rectangle 508 and the bucket emptying block 507 represents bits that
18 were taken from the bucket and allocated to the bucket emptying block 507.

19 When a block of a frame is simple, the “bucket” will fill, and when a block of the
20 frame is very complex the “bucket” will empty. All this is done by including as many
21 coefficients for a block as possible given a target size of bits in proportion to the non-zero
22 AC DCT coefficient bits in the corresponding block of the original encoded frame 501,

plus a share of any bits in the bucket. In general, the number (m) of non-zero AC DCT coefficients to keep for the block is found such that:

$$\text{ABS}(\text{Sum}(\text{Bits}(\text{DCT}(1))+\dots+\text{Bits}(\text{DCT}(m)) - \text{Breduced}(j)) < \text{Bucket}/\text{NBR}(j)$$

Where Bucket is the number of bits available in the “bucket” (size of the “bucket”), and NBR(j) is the remaining number of blocks in the frame including block (j), calculated as NB-j, so that the share of any bits in the bucket is Bucket/NBR(j). The DCT coefficients are obtained by parsing the original encoded block. The non-zero AC DCT coefficients that are retained in the reduced frame could be selected in the order of parsing, or in the order of the largest magnitude coefficients. In any case, the “bucket” remains as empty as possible considering that excess bits from one block will be shared with at least a certain number (NB) of following blocks.

The use of the bucket as a smoothing mechanism with a frame and between frames provides a significant improvement in picture quality because there is a large variability between the number of bits per DCT coefficient due to the MPEG-2 encoding code books and the number of escape sequences. For example, suppose Bucket = 36000 bits, NB=48*30=1440, SUM(DCT(1) + ... + DCT(11)) = 200, SUM(DCT(1) + ... + DCT(12)) = 300, and Breduced(j) = 280. The formula above will result in:

$$\text{ABS}(200-280) = 80 > 36000/1440 = 25$$

$$\text{ABS}(300-280) = 20 < 36000/1440 = 25$$

1 So we keep 12 DCT coefficients. In some cases there are only 11 DCT coefficients in the
2 example and there is no choice but to fill the "bucket" with 80 more bits.

3 FIG. 31 shows a graph of frame size (i.e., number of bits per frame) as a function
4 of frame number for various traces. A first trace 521 is for the original-quality MPEG
5 coded video, and a second trace 522 is for the reduced-quality MPEG coded video
6 produced from the original-quality MPEG coded video using the adaptive bit rate
7 reduction method as introduced above and as further described below. FIG. 31 also
8 shows a moving average 523 of the first trace 521, the average 524 of the first trace, and
9 the target frame size 524 which is the average of the second trace 522. It can be noticed
10 from FIG. 31 that the size of the reduced frames are proportional to the size in bits of the
11 original encoded file. In the case presented in FIG. 31 the rate reduction is done on all
12 the encoded frames including I, B and P type frames. It should be seen that the moving
13 average frame size over about 15 frames for the original-quality MPEG coded video is
14 very smooth and almost constant in time. This property is not required or used in any way
15 in the adaptive bit rate reduction method. For other applications, this property could be
16 exploited for example to ease calculations because the size of the GOP is kept constant
17 by the encoder and combines with the smoothness of the moving average to ensure that
18 the transport stream (TS) bit rate is constant in time (i.e., CBR). Because this property is
19 not required for the adaptive rate reduction method, the method may be performed with
20 the same accuracy for frame sequences that are not encoded for CBR. For example, the
21 method may be used upon I frames removed from original-quality MPEG coded video to
22 generate trick files for which there is no correlation between I frames separated by large
23 time intervals and making prediction of frame size impossible.

1 FIG. 32 illustrates such a case of independent I frames. FIG. 32 includes a first
2 trace 531 for the I frames from original-quality MPEG coded video, a second trace 532
3 which is a moving average of the first trace 531, a third trace 533 for the reduced-quality
4 I frames, and a fourth trace 534 which is the target frame size for the reduced-quality I
5 frames. It should be apparent that the moving average frame size in this case is far from
6 being constant and still the reduced rate frame sizes are proportional to the sizes of the
7 original encoded I frames.

8 FIGS. 33 to 35 show a flowchart of programming of a digital computer for
9 adaptive bit rate reduction of MPEG coded video as introduced above. For example, the
10 programming could be used in the MPEG scaling program (38 in FIG. 1) in the stream
11 server computer (25 in FIG. 1) in the video file server 24 in FIG. 1. In a first step 541 of
12 FIG. 33, the bucket is cleared. In step 542, the first video frame in the original-quality
13 MPEG clip is parsed down to the level of 8x8 DCT blocks. Then in step 543, the DCT
14 coefficient bit rate reduction factor is determined, as will be further described below with
15 reference to FIGS. 36 and 37. In step 544, the first 8x8 block is identified, and a block
16 index (J) is set to zero. In step 545, the first 8x8 block is parsed. If the block has no non-
17 zero AC DCT coefficients, as tested in step 546, then the block cannot be reduced, and
18 execution branches to step 547. If the end of the frame has not been reached, then
19 execution branches from step 547 to step 548, to get the next block and increment the
20 block counter (J), and to parse the next block in step 545. If the end of the frame is
21 reached, then execution continues to step 549. In step 549, if the end of the clip is
22 reached, then the adaptive bit rate reduction for the clip is finished. Otherwise, execution

1 continues from step 549 to step 550, to get the next frame, and parse the next frame in
2 step 542.

3 In step 546, if there are non-zero AC DCT coefficients in the current block, then
4 execution continues to step 561 in FIG. 34. In step 561, the original block size (BS) is
5 scaled by the reduction factor (RF) to compute a desired number of bits (DNB) for the
6 reduced block size. In other words, the DNB is computed as the product of RF and BS.
7 In step 562, the number of bits (NBB) in the bucket that are available for encoding non-
8 zero AC DCT coefficients for the reduced block is calculated by dividing the number of
9 bits in the bucket (BUK) by the number of remaining blocks in the frame, including the
10 current block, calculated as $(NB-J)$, and they are “borrowed” from the bucket by
11 decrementing BUK by NBB. NB is the number of blocks in the frame, and J is the index
12 of the current block. (Alternatively, NBB could be calculated as BUK/NBK , which
13 would more freely spread the bits in the bucket from one frame to the next, rather than
14 tending to deplete the bucket by the end of the frame.) Then in step 563, the number of
15 bits available (NBA) for encoding non-zero AC DCT coefficients for the reduced block
16 are calculated by adding the desired number of bits (DN) to the number of bits (NBB)
17 borrowed from the bucket.

18 In step 564, the first non-zero AC DCT coefficient is obtained in the parsing
19 order. This is for an implementation that selectively retains the low-order coefficients in
20 the parsing order, such as the low-pass scaling method in step 225 of FIG. 13 and in FIG.
21 14 as described above. Other scaling methods could be used instead. For example, the
22 preferred methods selectively retain or eliminate the non-zero coefficients. If there are
23 sufficient computational resources, for example, then largest magnitude scaling as in step

227 of FIG. 13 and in FIG. 15 could be used, or the approximate largest magnitude scaling as in step 229 of FIG. 13 and in FIGS. 18-19 could be used. If largest magnitude scaling were used, then the first coefficient obtained in step 564 would be the AC coefficient having the largest magnitude. If approximate largest magnitude scaling were used, then the first coefficient obtained in step 564 would be the AC coefficient having the approximate largest magnitude; in other words, it would be the first AC coefficient produced by the hash table sorting procedure of FIG. 19.

In step 565, the procedure determines the number of bits (NBC) for run-length coding for the first non-zero coefficient assuming that the first coefficient will be included in the reduced block. This is a very easy task for the low-pass scaling method in step 225 of FIG. 13, because the number of bits for the run-length coding of an AC coefficient in the reduced block will be the same as the run-length coding of the AC coefficient in the original block. For other scaling methods, the AC coefficient may have to be decoded and re-encoded to determine how many additional bits need to be used for run-length coding to include the AC coefficient in the reduced block. Moreover, for other scaling methods, for a second or subsequent AC coefficient in the block, the addition of a subsequent AC coefficient to the reduced block may change the run-length coding of prior AC coefficients as well as run-length coding for the additional coefficient. After step 565, execution continues to step 571 of FIG. 35.

In step 571 of FIG. 35, the number of bits (NBC) required for encoding and including the non-zero AC DCT coefficient in the reduced block is compared to the number of bit available (NBA) for encoding non-zero AC DCT coefficients in the block. If NBC is not less than or equal to NBA, then execution branches to step 572 to put the

1 remaining unallocated bits into the bucket. Execution loops back from step 572 to step
2 547 of FIG. 33.

3 In step 571 of FIG. 35, if NBC is less than or equal to NBA, then execution
4 continues to step 573. In step 573 the coefficient is transferred to the reduced-rate MPEG
5 coded video. Then in step 574, the number of bits available (NBA) to encode the non-
6 zero AC DCT coefficients of the block is decremented by the number of bits used (NBC)
7 in encoding the non-zero AC DCT coefficients of the block. In step 575, execution
8 branches to step 572 if there are no more non-zero AC DCT coefficients in the block.
9 Otherwise, execution continues from step 575 to step 576 in FIG. 34. In step 576, the
10 next non-zero AC DCT coefficient is obtained in the parsing order, for the
11 implementation of the low-pass scaling method. For other scaling methods, the next
12 largest-magnitude or approximate next largest magnitude non-zero AC DCT coefficient
13 would be obtained in step 576 instead of the next non-zero AC DCT coefficient in the
14 parsing order.

15 FIG. 36 shows a flowchart of a subroutine for determining the bit rate reduction
16 factor (RF) for a reduction from an MPEG source having a known constant bit rate. In a
17 first step 581, the offset rate (S), of bits in the original-quality MPEG source that are not
18 bits of the AC DCT coefficients, is determined. Then in step 582, the coefficient bit rate
19 reduction factor (RF) is computed from the known constant bit rate (BO) and padding
20 (PD) of the original-quality MPEG source, the offset rate (S), and the desired reduced bit
21 rate (BR) of the reduced-quality MPEG coded video. In particular, the coefficient bit rate
22 reduction factor is computed as $RF = (BR - S) / (BO - PD - S)$.

FIG. 37 shows a flowchart of a subroutine for determining the coefficient bit rate reduction factor (RF) for a reduction from an MPEG source having an unknown or variable bit rate. This subroutine, for example, is called in step 543 of FIG. 33. In a first step 591 of FIG. 37, the video frame size in bits (VS) is determined, for example from the parsing of the video frame in the previous step 542 in FIG. 33. Then in step 592 of FIG. 37, a moving average (VAVS) of the video frame size is determined over the last N frames. For example, $VAVS = (VS(I-N-1) + VS(I-N-2) + VS(I-N-3) + \dots + VS(I)) / N$, where VS(I) denotes the video frame size of Ith video frame in the original-quality MPEG source, and N is the number of frames over which the moving average is computed. N defines a "frame window," which is, for example, the number of video frames over 0.5 seconds worth of video data. In general, the moving average is the result of a low-pass digital filtering operation upon the video frame size. The low-pass digital filter is implemented by a subroutine, such as the following:

/* Compute Moving Average of up to N samples of VS */

/* J is initially set to N-1 */

/* K is initially set to N-1 */

/* SUM is initially set to zero */

IF (K<0) THEN GO TO CASE2

/* Handle the case at the beginning of the clip where */

/* N samples of VS are not yet available */

SUM ← SUM + VS

VAVS ← SUM/(N-J)

VSBUF(J) ← VS

```

1      J←J-1
2      RETURN
3      /* Handle the case where N samples of VS have been loaded */
4      /* into the buffer VSBUF */
5      CASE2: IF (K<0) THEN K←N-1
6      SUM←SUM+VS - VSBUF(K)
7      VAVS ← SUM/N
8      VSBUF(K) ← VS
9      K←K-1
10     RETURN

```

In step 593 of FIG. 37, a target average video frame size (VRAVS) is calculated from an accuracy rate control factor (AR), the desired reduced rate (BR) of the reduced-quality MPEG coded video, and the video frame rate (FR) according to $VRAVS = AR \cdot BR / FR$. The accuracy rate control factor (AR) is a factor substantially equal to 1 that can be adjusted on a frame-by-frame basis to control the bit rate of the reduced-quality MPEG coded video. The accuracy rate control factor (AR) could also be a constant set sufficiently less than 1 to account for variability in the moving average of the frame rate so that the bit rate of the reduced-quality MPEG coded video will not exceed the desired bit rate BR, and stuffing could later be inserted in the reduced-quality MPEG coded video to obtain precisely the desired bit rate BR.

In step 594, the number of bits (BS) in the frame that are not bits of the AC DCT coefficients is determined, for example, from the parsing of the video frame in the prior

1 step 542 of FIG. 33. In the last step 595 of FIG. 37, the coefficient bit rate reduction
2 factor (RF) is computed as $VRAVS/VAVS$.

3 FIG. 38 shows that a video buffer model could be used for adjusting the bit level
4 in the bucket when the adaptive bit rate reduction procedure 603 (as in FIGS. 33-37)
5 produces a reduced quality MPEG-2 stream 604 in real time from a source 605 of
6 original-quality MPEG-2 coded video. The video buffer model 601 accumulates the
7 video buffer level (VBV) in accordance with the time stamps (DTS, PTS) and program
8 clock reference (PCR) values in the reduced-quality MPEG-2 stream 604 indicating when
9 bits from the stream 604 are added to the video buffer level (VBV) and when bit from the
10 steam 604 are removed from the video buffer level (VBV). In particular, the number of
11 bits in the bucket 602 is incremented when the video buffer level becomes less than a low
12 threshold TH1, and the number of bits in the bucket 602 is decremented when the video
13 buffer level becomes greater than a high threshold TH2. The thresholds TH1 and TH2
14 are selected to prevent underflow of the video buffer level (i.e., VBV becoming less than
15 zero) and to prevent overflow of the video buffer level (i.e., VBV exceeding a maximum
16 size that is greater than TH2). The video buffer model could be implemented as a
17 subroutine that would be called when a frame is encoded and also at the time when the
18 frame would be decoded, relative to the time that the frame is encoded. In this case, for
19 example, the maximum size of the buffer includes a multiplicity of the reduced size
20 frames, the low threshold is larger than the largest reduced frame, and the difference
21 between the size of the buffer and the high threshold is larger than the size of the largest
22 reduced frame. The subroutine could be coded as follows:

```

1      /* Video buffer model */
2
3      /* This routine is called when a reduced frame has been encoded. */
4
5      /* Before increasing the video buffer level for the frame just encoded, */
6
7      /* a video buffer queue (VBUFQUEUE) is inspected. The video buffer */
8
9      /* queue has respective entries for frames that have been encoded */
10
11     /* but not yet decoded. Each entry includes a decode time stamp */
12
13     /* indicating when the frame is to be removed from the video buffer */
14
15     /* relative to the loading of the current frame into the video buffer. */
16
17     START: IF (VBUFQUEUE IS EMPTY) THEN GO TO STEP2
18
19     STEP1: GET DECODE TIME STAMP OF ENTRY AT HEAD OF VBUFQUEUE
20
21           IF(DECODE TIME STAMP IS AFTER LOADING TIME OF CURRENT
22
23           FRAME) THEN GO TO STEP 2
24
25           GET FRAME SIZE (FSHE) OF FRAME OF ENTRY AT HEAD OF
26
27           VBUFQUEUE
28
29           VBUF←VBUF-FSHE
30
31           REMOVE ENTRY FROM HEAD OF VBUFQUEUE
32
33           GO TO START
34
35     STEP2: IF (VBUF < TH1) THEN BUK←BUK-(VBUF-TH1)
36
37           IF (VBUF < 0) THEN (VBUF ←0; REPORT UNDERFLOW ERROR)
38
39           VBV←VBV+FS
40
41           IF (VBV) > TH2 THEN BUK←BUK-(VBV-TH2)
42
43           IF (VBV>MAX) THEN (VBV ←MAX; REPORT OVERFLOW ERROR)
44
45           PUT FS and relative time stamp of frame on buffer queue

```

1 RETURN

2
3 Although the adaptive bit rate reduction method has been shown in FIG. 38 being
4 used for the real-time production of a reduced-quality MPEG-2 stream 604 from a source
5 605 of original-quality MPEG-2 coded video, the same method can be used in order to
6 reduce the size of an MPEG-2 transport stream file to a given file size either for storage
7 on CD ROMs or for trick file generation.

8 For trick file generation, sometimes there is need to insert freeze frames to
9 compensate for large variations in the frame size. For generating trick files for video-on-
10 demand (VOD) applications, in particular, the file size is generally limited to 10% of the
11 original file size while the target reduced bit rate must be smaller or equal than the bit
12 rate of the original file. When only I frames are used the equivalent bit rate of the I
13 frames of an IBBP encoded clip is 3 to 4 times higher than the bit rate of the encoded file,
14 for example for clips with lots of action encoded at 4 Mbps by a high quality encoder, the
15 equivalent bit rate of the I frames of the clip is 14.5 Mbps. Therefore, the bit rate
16 reduction factor is 3.6, which at such low original bit rate will generate bad quality video
17 frames. To compensate for this large reduction, some of the I frames in the original file
18 are discarded and replaced by P freeze frame that are very small in size. The selection of
19 I frames to discard is done based on the number of bits that will be allocated per block at
20 such a low bit rate as 4 Mbps.

21 In view of the above, there has been described a method of producing reduced bit
22 rate MPEG coded video while achieving uniform proportional video quality reduction
23 proportional to the ratio of the original frame size and the scaled down frame size by

1 proportional reduction of block size and smoothing between frames using a “bucket” of
2 bits. Bits are placed in the bucket when the bits available to be allocated to a block
3 cannot be allocated to the block because there is a rather coarse granularity in the number
4 of bits that can be allocated to a block. A certain fraction of the bits in the bucket are
5 made available for encoding AC DCT coefficients in the block, so that the bits in the
6 bucket after allocation to a block are available for encoding AC DCT coefficients in at
7 least a certain number of following blocks. The method can be applied to a source of
8 original-quality MPEG coded video having an unknown or variable bit rate by computing
9 a moving average of the number of bits per frame over a certain number of consecutive
10 frames. The method improves the quality of the reduced-quality MPEG coded video for
11 a given bit rate by increasing the bit allocation to the video by reducing to a minimum the
12 amount of padding used to precisely achieve the desired bit rate, and by obtaining a
13 substantially uniform degradation in quality over the frames and blocks in the frames.
14 The method can be done in the compressed domain without decoding of the run-length
15 encoded AC DCT coefficients, so that it can be done in real time while streaming MPEG-
16 2 data from a file server at a variable reduced rate to accommodate network congestion.
17 The method can also be used to produce a trick file in parallel with the encoding of an
18 MPEG-2 stream for storage of compressed video in a video file server.

19

20